

HP ProLiant Storage Server administration guide

Part number: Part Number: 378127-002
Second Edition edition: (March 2005)



Legal and notice information

Copyright © 2005 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft®, MS Windows®, Windows®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

HP ProLiant Storage Server administration guide

Contents

About this Guide	13
Intended audience	13
Prerequisites	13
Conventions	13
Document conventions	13
Text symbols	13
Getting help	14
HP technical support	14
HP storage web site	15
HP authorized reseller	15
1 System Overview	17
Product definition and information	17
Server hardware and software features	17
Product information	17
Product manageability	17
Product redundancy	18
Deployment scenarios	18
Environment scenarios	19
Workgroup	20
Domain	20
User interfaces	20
Storage server web-based user interface	20
Storage server desktop	21
Storage Server Management Console	22
NIC Team Setup	22
2 Basic Administrative Procedures and Setup Completion	23
Basic administrative procedures	23
Setting the system date and time	24
Shutting down or restarting the server	24
Viewing and maintaining audit logs	25
Using Remote Desktop	154
Improper closure of Remote Desktop	27
Setting up e-mail alerts	27
Changing system network settings	28
Setup completion	29
Managing system storage	29
Creating and managing users and groups	29
Creating and managing file shares	30
Activating the iLO port using the license key	30
Setting up Ethernet NIC Teams (Optional)	30
3 Disk and Volume Management	31
Storage servers with configurable storage	31
Storage configuration overview	32
Step 1: Create disk arrays	33
Step 2: Create logical disks from the array space	33
Step 3: Verify newly created logical disks	33

Step 4: Create a volume on the new logical disk	33
Array Configuration Utility (Smart Array-based storage only)	34
Using the ACU to configure storage	34
ACU guidelines	37
Managing disks on configurable storage servers	37
Creating a new volume via the WebUI	38
Advanced Disk Management	40
Guidelines for managing disks	41
Volumes page	52
Managing volumes	43
Managing disks after quick restore	48
Storage servers with pre-configured storage	48
Disk Management utility	49
Disk Management guidelines	50
Adaptec Storage Manager	51
Volumes page	52
Scheduling defragmentation	53
Disk quotas	55
Enabling quota management	55
Setting user quota entries	56
DiskPart	58
Example of using DiskPart	59

4 Shadow Copies	61
Overview	79
Shadow copy planning	61
Identifying the volume	62
Allocating disk space	62
Converting basic storage disks to dynamic disks	63
Identifying the storage area	63
Determining creation frequency	64
Shadow copies and drive defragmentation	64
Mounted drives	64
Managing shadow copies	65
The shadow copy cache file	67
Enabling and creating shadow copies	68
Viewing a list of shadow copies	69
Set schedules	69
Scheduling shadow copies	69
Deleting a shadow copy schedule	69
Viewing shadow copy properties	70
Disabling shadow copies	71
Managing shadow copies from the storage server desktop	72
Shadow Copies for Shared Folders	72
SMB shadow copies	73
NFS shadow copies	74
Recovery of files or folders	75
Recovering a deleted file or folder	75
Recovering an overwritten or corrupted file	76
Recovering a folder	76
Backup and shadow copies	77

5 User and Group Management	79
Overview	79
Domain compared to workgroup environments	79
User and group name planning	79
Managing user names	79
Managing group names	80
Workgroup user and group management	80

Managing local users	80
Adding a new user	81
Deleting a user	82
Modifying a user password	82
Modifying user properties	82
Managing local groups	83
Adding a new group	83
Deleting a group	84
Modifying group properties	84
6 Folder, Printer, and Share Management	87
Folder management	87
Navigating to a specific volume or folder	87
Creating a new folder	88
Deleting a folder	89
Modifying folder properties	89
Creating a new share for a volume or folder	90
Managing shares for a volume or folder	91
Managing file level permissions	92
Share management	98
Share considerations	99
Defining Access Control Lists	99
Integrating local file system security into Windows domain environments	99
Comparing administrative (hidden) and standard shares	99
Planning for compatibility between file sharing protocols	100
NFS compatibility issues	100
Managing shares	100
Creating a new share	133
Deleting a share	135
Modifying share properties	135
Protocol parameter settings	106
DFS protocol settings	107
Deploying DFS	107
DFS Administration Tool	108
Accessing the DFS namespace from other computers	109
Setting DFS sharing defaults	109
Creating a local DFS root	110
Deleting a local DFS root	111
Publishing a new share in DFS	111
Publishing an existing share in DFS	113
Removing a published share from DFS	113
Storage management	114
Directory quotas	115
Establishing directory quotas	116
File screening	117
Storage reports	118
Print services (where licensed)	119
Configuring the print server	119
Removing the print server role	121
Adding an additional printer	121
Adding additional operating system support	123
Installing print services for UNIX	123
HP Web Jetadmin	123
7 Services for NFS/UNIX	125
Server for NFS	125
Authenticating user access	125
S4U2 functionality	126
Indicating the computer to use for the NFS user mapping server	127

Logging events	128
Server for NFS server settings	129
Installing NFS Authentication software on the domain controllers and Active Directory domain controllers	130
Installing SFU 3.5 from CD	131
Understanding NTFS and UNIX permissions	132
NFS file shares	133
Creating a new share	133
Deleting a share	135
Modifying share properties	135
Anonymous access to an NFS share	137
NFS only	138
NFS protocol properties settings	138
NFS async/sync settings	139
NFS locks	140
NFS client groups	141
Adding a new client group	142
Deleting a client group	143
Editing client group information	143
NFS user and group mappings	144
Types of mappings	145
Explicit mappings	145
Simple mappings	145
Squashed mappings	145
User name mapping best practices	146
Creating and managing user and group mappings	147
General tab	148
Simple mapping tab	148
Explicit user mapping tab	149
Explicit group mapping tab	150
Backing up and restoring mappings	152
Backing up user mappings	152
Restoring user mappings	153
Creating a sample NFS file share	153
Remote Access	154
Using Remote Desktop	154
Using Telnet Server	155
Using Remote Shell Service	155
Interix	155
Shells	156
Programming Languages	156
Enabling setuid behavior for Interix programs	156

8 NetWare File System Management 157

Installing Services for NetWare	157
Managing File and Print Services for NetWare	158
Creating and managing NetWare users	160
Adding local NetWare users	160
Enabling local NetWare user accounts	160
Managing NCP volumes (shares)	161
Creating a new NCP share	162
Modifying NCP share properties	164

9 Remote Access Methods and Monitoring 165

Web-based user interface	165
Remote Desktop	165
Telnet Server	165
Enabling Telnet Server	166
Sessions information	166

Integrated Lights-Out port	166
Features	167
Security features	167
Manage Users feature	167
Manage Alerts feature	167
Integrated Lights-Out Port configuration	168
Using the Integrated Lights-Out Port to Access the Storage Server	168
HP Insight Manager Version 7	169
10 Cluster Administration	171
Cluster overview	171
Multi-node support beyond two nodes	171
Cluster terms and components	172
Nodes	172
Resources	172
Virtual servers	172
Failover	173
Quorum disk	173
Cluster concepts	173
Sequence of events for cluster resources	173
Hierarchy of cluster resource components	174
Cluster planning	175
Storage planning	176
Network planning	176
Protocol planning	177
Preparing for cluster installation	178
Before beginning installation	178
Using Secure Path	178
Uninstalling Storage Manager	179
Checklists for cluster server installation	179
Network requirements	180
Shared disk requirements	180
Cluster installation	180
Setting up networks	181
Configuring the private network adapter	181
Configuring the public network adapter	181
Renaming the Local Area Network icons	182
Verifying connectivity and name resolution	182
Verifying domain membership	182
Setting up a cluster user account	182
About the Quorum disk	182
Configuring shared disks	183
Verifying disk access and functionality	183
Configuring cluster service software	183
Creating a cluster	183
Adding nodes to a cluster	185
Geographically dispersed clusters	186
HP ProLiant Storage Server software updates	186
Cluster groups and resources, including file shares	187
Cluster group overview	187
Node-based cluster groups	187
Load balancing	211
Cluster resource overview	188
File share resource planning issues	188
Resource planning	188
Permissions and access rights on share resources	189
NFS cluster-specific issues	189
Non cluster aware file sharing protocols	190
Creating a new cluster group	190
Adding new storage to a cluster	191

Creating physical disk resources	192
Creating file share resources	193
Setting permissions for an SMB file share	194
Creating NFS share resources	196
Setting permissions for an NFS share	197
Creating IP address resources	198
Creating network name resources	199
Basic cluster administration procedures	200
Failing over and failing back	200
Restarting one cluster node	201
Shutting down one cluster node	201
Powering down the cluster	201
Powering up the cluster	202
Shadow copies in a clustered environment	202
Creating a cluster printer spooler	203
A NIC Teaming	205
Installing the HP Network Teaming Utility	205
Opening the HP Network Teaming Utility	207
Adding and configuring NICs in a team	207
Fault tolerance	210
Load balancing	211
Configuring the NIC team properties	212
Renaming the teamed connection	212
Showing a connection icon on the taskbar	212
Configuring the TCP/IP protocol on the new team	212
Checking the status of the team	214
NIC teaming troubleshooting	215
Index	217

Figures

1 Storage server desktop	22
2 Maintenance tab	24
3 Date and Time page	24
4 Shutdown page	25
5 Logs page	25
6 Remote Desktop session	27
7 Network tab	29
8 Disks menu—configurable storage models	32
9 Array Management page	35
10 Systems Management Homepage	36
11 Manage Disks page—configurable storage server	38
12 Creating a new volume, page 1	39
13 Creating a new volume, page 2	40
14 Disk Management utility	50
15 Volumes page	88
16 Manage Volumes page	44
17 Expanding a LUN (Smart Array only)	45
18 Extending a volume (basic disk)	46
19 Extending a volume (dynamic disk)	47
20 Disks tab—medium and small business class	49
21 Disk Management utility	50
22 Adaptec Storage Manager	88
23 Volumes tab	53
24 Setting user quotas	57
25 Add new quota entry	57
26 Shadow Copies page	66
27 Shadow copies stored on source volume	67
28 Shadow copies stored on separate volume	67
29 Shadow Copy Properties page	71
30 Accessing Shadow Copies from My Computer	72
31 Client GUI	74
32 Recovering a deleted file or folder	76
33 Local Users page	81
34 Create New User page	82
35 User Properties page	83
36 Local Groups page	83
37 Create New Group page, General tab	84
38 Group Properties page, General tab	85
39 Group Properties page, Members tab	86
40 Volumes page	88
41 Folders page	88
42 Create a New Folder page, General tab	89
43 Folder Properties page, General tab	90
44 Create New Share page, General tab	134
45 Properties dialog box, Security tab	93
46 Advanced Security Settings dialog box, Permissions tab	94
47 User or Group Permission Entry dialog box	95
48 Advanced Security Settings dialog box, Auditing tab	95
49 Select User or Group dialog box	96
50 Auditing Entry dialog box for folder name NTSF Test	97
51 Advanced Security Settings dialog box, Owner tab	98
52 Create a New Share page, General tab	134
53 Share Properties page, General tab	135
54 Share Properties page, Windows Sharing tab	103

55 Share Properties page, UNIX Sharing tab	104
56 Local Area Connection Properties page, Install option	105
57 File Sharing Protocols page	107
58 DFS Win32 GUI	108
59 DFS Properties page, General tab	110
60 DFS Properties page, Local DFS Root tab	111
61 DFS share example	112
62 DFS share example, mapped drive	113
63 Uninstall storage manager	115
64 Directory Quota Policies page	117
65 Microsoft Services for NFS screen, Settings tab	128
66 Server for NFS screen, Logging tab	129
67 Server for NFS screen, Server Settings tab	130
68 Create a New Share page, General tab	134
69 Share Properties page, General tab	135
70 UNIX Sharing tab	136
71 NFS Sharing Protocols page	139
72 NFS Async/Sync Settings page	140
73 NFS Locks page	141
74 NFS Client Groups page	142
75 New NFS Client Group page	143
76 Edit NFS Client Groups page	144
77 Mapping server "ls -al" command example	146
78 User and Group Mappings page, General tab	148
79 User and Group Mappings page, Simple Mapping tab	149
80 User and Group Mappings page, Explicit User Mapping tab	150
81 User and Group Mappings page, Explicit Group Mapping tab	151
82 User Name Mapping screen, Map Maintenance tab	152
83 Local Area Connection Properties page, Install option	158
84 Installing File and Print Services for NetWare	158
85 File and Print Services for NetWare dialog box	159
86 New User dialog box	160
87 NetWare Services tab	161
88 Create Volume dialog box	162
89 Access Through Share Permissions dialog box	163
90 Add Users and Groups dialog box	163
91 Storage server cluster diagram	172
92 Cluster concepts diagram	174
93 Uninstall Storage Manager	179
94 Cluster tab	183
95 Adding a new node	185
96 Cluster update tool	186
97 Cluster Groups page	191
98 Cluster Resources page	192
99 Creating a file share resource	194
100 Resource parameters for SMB file share	195
101 Set resource permissions	196
102 NFS Share Resource parameters	197
103 Set NFS Share resource permissions	198
104 Creating an IP address resource	199
105 Network Name Parameters	200
106 Installing Network Teaming	206
107 Network Teaming installation complete	206
108 HP Network Teaming utility icon	207
109 HP Network Teaming Properties dialog box	208
110 NIC Properties, Teaming Controls tab, Fault Tolerant option	209
111 HP Network Teaming dialog box	210
112 NIC Properties, Teaming Controls tab, Load Balancing option	211
113 NIC Team Properties dialog box	213
114 NIC Team TCP/IP Properties dialog box	214

115 NIC Teaming status	215
----------------------------------	-----

Tables

1 Document conventions	13
2 WebUI main menu tabs	21
3 Welcome screen contents	21
4 Disks tab options	32
5 Manage Disks options	38
6 Volumes page object/task selector	43
7 Manage Volumes options	44
8 Disks tab options	49
9 Volumes page object/task selector	53
10 Common DiskPart commands	59
11 Shadow Copies fields	66
12 Shadow Copies tasks	66
13 Group name examples	80
14 Command Line Interface Command Prompts	154
15 Sharing protocol cluster support	178
16 Power sequencing for cluster installation	181
17 NIC Teaming Troubleshooting	215

About this Guide

Intended audience

This book is intended for system administrators who are experienced with setting up and managing a network server.

Prerequisites

Before beginning, make sure you consider the items below.

- Knowledge of the Microsoft® Windows® Storage Server 2003 operating system
- Knowledge of HP hardware
- Location of all documentation shipped with your server

Conventions

Conventions consist of the following:

- Document conventions
- Text symbols

Document conventions

This document follows the conventions in [Table 1](#).

Table 1 Document conventions

Convention	Element
Blue text: Figure 1	Cross-reference links
Bold	Menu items, buttons, keys, tabs, and box names
<i>Italics</i>	Text emphasis and document titles in body text
Monospace font	User input, commands, code, file and directory names, and system responses (output and messages)
<i>Monospace, italic font</i>	Command-line and code variables
Blue underlined sans serif font text (http://www.hp.com)	Web site addresses

Text symbols

The following symbols may be found in the text of this guide. They have the following meanings:



WARNING:

Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or death.



CAUTION:

Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.



IMPORTANT:

Text set off in this manner provides additional help to readers by providing essential information to help avoid problems.



NOTE:

Text set off in this manner presents commentary, sidelights, or interesting points of information.

Getting help

If you still have a question after reading this guide, contact an HP authorized service provider or access our web site: <http://www.hp.com>.

HP technical support

Telephone numbers for worldwide technical support are listed on the following HP web site: <http://www.hp.com/support/>. From this web site, select the country of origin.



NOTE:

For continuous quality improvement, calls may be recorded or monitored.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

HP storage web site

The HP web site has the latest information on this product, as well as the latest drivers. Access storage at: <http://www.hp.com/country/us/eng/prodserv/storage.html>. From this web site, select the appropriate product or solution.

HP authorized reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-282-6672
- In Canada, call 1-800-863-6594
- Elsewhere, see the HP web site for locations and telephone numbers: <http://www.hp.com>.

1 System Overview

The HP ProLiant Storage Server products can be used in many types of computing environments, from basic Microsoft Windows workgroups to complicated multiprotocol domains using DFS, NFS, FTP, HTTP, and Microsoft SMB. The corresponding varieties of clients that can be serviced include any Windows, UNIX, Linux, Novell, or Macintosh variant.

This chapter provides an overview of these environments and deployments and includes a brief descriptions of the available user interfaces.

Product definition and information

The HP ProLiant Storage Server family of products includes enterprise class, as well as remote office or small to medium business class solutions that provide reliable performance, manageability, and fault tolerance.



NOTE:

The HP ProLiant Storage Server Installation Guide includes a chart to help users determine which features discussed in this administration guide apply to a specific model.

Server hardware and software features

Refer to the HP ProLiant Storage Server QuickSpecs for a list of server hardware and software features available on the HP web site: <http://www.hp.com/go/proliant>

Product information

The storage server provides performance gains over general purpose servers by integrating optimized hardware components and specialized software. Integrating storage servers into the network improves the performance of existing servers because storage servers are optimized for file serving tasks.



NOTE:

Each HP ProLiant Storage Server has been specifically designed to function as a network attached storage server. Unless specifically authorized by HP, do not use the server software to support additional applications or significant functionality other than system utilities or server resource management, or similar software that you install and use solely for system administration, system performance enhancement, and/or preventative maintenance of the server.

Product manageability

The storage server ships with the following utilities and features:

- **Rapid Startup Wizard**—User friendly configuration utility that ensures easy configuration.
- **Web-based user interface (WebUI)**—Simple, graphical user interface that helps with administration tasks.
- **Ability to connect directly to the console.**
- **Insight Manager** (not available on all models)—Monitors the operations of HP servers, workstations, and clients. Insight Manager provides system administrators more control through comprehensive fault and configuration management, and remote management.
- **Integrated Lights-Out feature** (not available on all models)—Provides remote access, sends alerts, and performs other management functions, even if the operating system of the host server is not responding.

Product redundancy

The storage server is specifically designed to perform file serving tasks for networks, using industry standard components to ensure reliability.

Other industry standard features, such as redundant array of independent drives (RAID) and remote manageability, further enhance the overall dependability of the storage server.

The clustering ability of select storage servers further ensures continuous data availability because data being processed by one server transitions over to the other server in a failover situation.

Deployment scenarios

Various deployment scenarios are possible. See the HP ProLiant Storage Server installation guide for configurations. Typical application of storage servers include:

- **File server consolidation**

As businesses continue to expand their information technology (IT) infrastructures, they must find ways to manage larger environments without a corresponding increase in IT staff. Consolidating many servers into a single storage server decreases the number of points of administration, and increases the availability and flexibility of storage space.

- **Multiprotocol environments**

Some businesses require several types of computing systems to accomplish various tasks. The multiprotocol support of the storage server allows it to support many types of client computers concurrently.

- **Protocol and platform transitions**

When a transition between platforms is being planned, the ability of the storage server to support most file sharing protocols allows companies to continue to invest in file storage space without concerns about obsolescence. For example, an administrator planning a future transition from Windows to Linux can deploy the storage server with confidence that it can support both CIFS and NFS simultaneously, assuring not only a smooth transition, but also a firm protection of their investment.

- **Remote office deployment**

Frequently, branch offices and other remote locations lack dedicated IT staff members. An administrator located in a central location can use the WebUI of the storage server, Microsoft Terminal Services, and other remote administration methods to configure and administer all aspects of the storage server.

- **Microsoft Windows Storage Server 2003 Feature Pack deployment**

The Feature Pack is available for select storage servers. The Feature Pack allows Microsoft Exchange Server 2003 databases and transaction logs to be stored on a storage server running Microsoft Windows Storage Server 2003. A single storage server computer running the Feature Pack can host the databases and transaction logs of up to two Exchange servers and up to 1,500 Exchange mailboxes.

The Feature Pack installs new components on both the storage server computer and Exchange Server 2003. These components provide tools and services that allow Exchange databases and transaction logs to be moved to a storage server computer, and they perform the necessary configuration updates to give Exchange Server 2003 access to the remotely stored files.

Environment scenarios

The storage server is deployed in one of two security modes:

- Workgroup
- Domain (Windows NT® Domain or Active Directory Domain)

The storage server uses standard Windows user and group administration methods in each of these environments.

Regardless of the deployment, the storage server integrates easily into multiprotocol environments, supporting a wide variety of clients. The following protocols are supported:

- Distributed File System (DFS)
- Network File System (NFS)
- Hypertext Transfer Protocol (HTTP)
- File Transfer Protocol (FTP)
- Microsoft Server Message Block (SMB)

Workgroup

In a workgroup environment, users and groups are stored and managed separately, on each member server of the workgroup. Workgroups are typical for very small deployments where little or no computing environment planning is required.



NOTE:

In a clustered deployment (servers only), the clusters must be members of a domain. Therefore, workgroup environments are only supported in non-clustered deployments.

Domain

When operating in a Windows NT or Active Directory domain environment, the storage server is a member of the domain and the domain controller is the repository of all account information. Client machines are also members of the domain and users log on to the domain through their Windows-based client machines. The domain controller also administers user accounts and appropriate access levels to resources that are a part of the domain. Additional information about planning for domain environments can be found at:

<http://www.microsoft.com/windowsserver2003/technologies/directory/activedirectory/default.mspx>

The storage server obtains user account information from the domain controller when deployed in a domain environment. The storage server itself cannot act as a domain controller, backup domain controller, or the root of an Active Directory tree as these functions are disabled in the operating system.

User interfaces

There are several user interfaces that administrators can use to access and manage the storage server. Two of these interfaces are:

- Storage server WebUI
- Storage server desktop

Each interface contains the same or similar capabilities, but presents them in a different manner. Each of these interfaces are illustrated in the following sections.

Storage server web-based user interface

The WebUI provides system administration, including user and group management, share management, and local storage management.

Refer to the *HP ProLiant Storage Server Installation Guide* for detailed information on using the Rapid Startup Wizard for initial setup.

To access the WebUI, launch a web browser, and then enter the following in the address field:

`https://<your machine name or IP Address>:3202/`

The default user name is Administrator. The default password is hpinvent. Online help for the WebUI is available by clicking the **Help** tab on the primary WebUI screen.

Table 2 WebUI main menu tabs

Tab	Description
Status	View alerts generated by the WebUI.
Network	Access system settings, including system identification, global settings, interfaces settings, administration settings, Telnet settings, and SNMP settings.
Disks	Manage disks, volumes, disk quotas, and shadow copies.
Users	Manage local users and groups.
Shares	Create folders and shares to control access to files. Define protocols and protocol parameters.
Array Management (select models)	Manage arrays and pathing software.
Maintenance	Access maintenance tasks including setting date and time, performing system restarts and shutdowns, viewing audit logs, setting up e-mail alerts, linking to remote management, and selecting and configuring the UPS.
HP Utilities	Access HP system management utilities such as File and Print Services for NetWare.
Cluster (select models)	Configure and manage clusters.
Help	Access help information for the WebUI.

Table 3 Welcome screen contents

Tab	Description
Installation Overview (not on all models)	Set up and configure the storage server. This is an online, supplemental guide. A more comprehensive paper document is provided in the country kit that shipped with the server.
Take a Tour	Learn how to use the storage server.
Rapid Startup Wizard	Enter system setup and configuration information.
Set Administrator Password	Create a password for the storage server administrator.
Set Server Name	Choose a name so that client computers can connect to the server.
Set Default Page	Choose which page the storage server displays first.

Storage server desktop

The storage server desktop can be accessed by:

- Directly connecting a keyboard, mouse, and monitor.
- Using the WebUI **Maintenance** tab, and then selecting **Remote Desktop**.
- Using the Integrated Lights-out port (not available on all models).



NOTE:

When using Remote Desktop to connect to the storage server desktop do not use the window close feature (X). Select Start > Log Off Administrator to exit Remote Desktop.




Figure 1 Storage server desktop

The following icons are available from the desktop:

- Storage Server Management Console
- NIC Team Setup

Storage Server Management Console

Click this icon to access the following folders:

- **Core Operating System**—Used to manage local users and groups, access performance logs and alerts, and manage the event viewer.
- **Disk System**—Contains access to the HP Array Configuration Utility (select models), and local disk management, including a volume list and a graphical view of the disks.
- **File Sharing**—Contains modules for the configuration of file sharing exports. CIFS/SMB (Windows) and NFS (UNIX) file shares are managed through this folder.
- **System**—Contains system summary information.

NIC Team Setup

Click the **NIC Team Setup** icon to install the HP Network Teaming and Configuration utility. See Appendix A for additional information on this feature.



NOTE:

The HP Network Teaming and Configuration utility is not supported nor available on all models.

2 Basic Administrative Procedures and Setup Completion

Basic system administration functions are discussed in this chapter.

This chapter also continues the process of setting up the system that was started using the *HP ProLiant Storage Server Installation Guide* by discussing additional setup procedures and options.

Unless otherwise instructed, all procedures are performed using the storage server web-based user interface (WebUI).



NOTE:

The storage server desktop can be accessed via a directly connected keyboard, mouse, and monitor or through Remote Desktop. Select models can use a RILOE or iLO port.

Basic administrative procedures

Basic administrative procedures include:

- Setting the system date and time
- Shutting down or restarting the server
- Viewing and maintaining audit logs
- Using Remote Desktop
- Setting up e-mail alerts
- Changing system network settings

These functions are performed in the **Maintenance** tab of the WebUI except for changing system network settings, which is in the **Network** tab.




Figure 2 Maintenance tab

Setting the system date and time

To change the system date or time:

1. From the WebUI, click **Maintenance**, and then **Date/Time**. The **Date and Time Settings** page is displayed.
2. Enter the new values, and then click **OK**.

Figure 3 Date and Time page

Shutting down or restarting the server



CAUTION:

Notify users before powering down the system. Both UNIX and Windows NT users can be drastically affected if they are not prepared for a system power-down.

1. From the storage server WebUI, click **Maintenance, Shutdown**. Several options are displayed: **Restart**, **Shut Down**, and **Scheduled Shutdown**.

Maintenance

Date/Time | Shutdown | Logs | Backup | Remote Desktop | Alert E-Mail | Language | UPS

Shutdown
Shut down or restart the server immediately or at a scheduled time.

- **Restart**
Immediately shut down and then automatically restart the server.
- **Shut Down**
Immediately shut down and power off the server.
- **Scheduled Shutdown**
Schedule a shutdown or restart to occur later.

Shutdown Related Alerts	Shutdown Related Status
	Shared Files 0 files open

Figure 4 Shutdown page

- a. To shut down and automatically restart the server, click **Restart**.
- b. To shut down and power off the server, click **Shut Down**.
- c. To schedule a shutdown, click **Scheduled Shutdown**.

2. Regardless of the choice, a confirmation prompt is displayed. After verifying that this is the desired action, click **OK**.



NOTE:

Client computers do not receive a warning message prior to shutdown.

Viewing and maintaining audit logs

A variety of audit logs are provided on the storage server. System events are grouped into similar categories, representing the seven different logs.

To access the logs from the WebUI, click **Maintenance, Logs**. The **Logs** page is displayed.

Maintenance

Date/Time | Shutdown | **Logs** | Backup | Remote Desktop | Alert E-Mail | Language | UPS

Logs
View, clear, download, and configure logs.

- **Application Log**
The application log contains events logged by programs.
- **System Log**
The system log contains events logged by the operating system components.
- **Security Log**
The security log can record security events such as valid and invalid logon attempts.
- **Web Administration Log**
The web administration log contains events logged by the Web server related to access to the administration web site.
- **NFS Log**
The NFS log contains events logged by the NFS server.
- **FTP Log**
The FTP log contains events logged by the FTP server.
- **Web (HTTP) Shares Log**
The Web (HTTP) Shares log contains events logged by the Web server related to accessing the HTTP Shares.

Back

Figure 5 Logs page

A variety of logs are available and are listed in [Figure 5](#).

Each log has viewing, clearing, printing, and saving options.



NOTE:

You should not use the WebUI to view log files greater than 2 MB. Select Log properties to adjust the maximum file size, or download the file to view.



NOTE:

NFS logging is disabled by default. Enable NFS logging using the Management Console. NFS stops logging when the log file is full.

Using Remote Desktop

Remote Desktop is provided in the WebUI to allow for additional remote system administration and the use of approved third-party applications. Backup software and antivirus programs are examples of approved applications.

To open a Remote Desktop session from the WebUI, select **Maintenance, Remote Desktop**. A Remote Desktop session is opened. Enter the appropriate password to log on to the server.




Figure 6 Remote Desktop session



CAUTION:

Two open sessions of Remote Desktop are allowed to operate at the same time. After completing an application do not use the window close feature () to close that session of Remote Desktop. Select **Start > Log Off Administrator** to exit Remote Desktop.

Improper closure of Remote Desktop

Certain operations can leave the utilities running if the browser is closed versus exiting from the program via the application menu or logging off the Remote Desktop session. A maximum of two Remote Desktop sessions may be used at any given time. Improper exit from a session can result in the sessions becoming consumed. Sessions and processes can be terminated using the **Terminal Services Manager** via **Start > Programs > Administrator Tools**.



NOTE:

The Terminal Services Manager must be accessed via the direct attached console, or via the RILOE or iLO port on select models.

Setting up e-mail alerts

E-mail messages are limited to the alerts generated from the WebUI status bar or the WebUI status page, as well as some event log messages. Some alerts, such as the restart of the server, only occur if the WebUI

was utilized to initiate the request. For example, a restart initiated using the WebUI generates an e-mail message indicating a restart has occurred. Initiating a restart using the Windows Storage Server 2003 schedule or Desktop will not. Messaging in the status bar and page is limited to the following areas:

- WebUI Alerts
 - NTBackup backup started
 - NTBackup restore started
 - Defrag started
 - UPS power failure
 - Restart pending
 - Shutdown pending
 - DFS not configured
 - Date and time not configured
 - No certificate
 - Quota management alerts
- Event Log Messages
 - NTBackup Information
 - UPS power failed
 - UPS power restored
 - UPS invalid config
 - UPS system shutdown
 - Quota management alerts

To activate this option:

1. From the WebUI, click **Maintenance**. Then click **Alert E-mail**. The **Set Alert E-Mail** page is displayed.
2. Select **Enable Alert E-mail**.
3. Indicate the types of messages to be sent.
 - Critical alerts
 - Warning alerts
 - Informational alerts
4. Enter the desired e-mail address in the appropriate boxes.
5. After all settings have been entered, click **OK**.

Changing system network settings

Network properties are entered and managed from the **Network** tab. Most of these settings are entered as part of the Rapid Startup process. Settings made from this tab include adding the storage server to a domain.

Online help is available for these settings. [Figure 7](#) is an illustration of the **Network** tab.




Figure 7 Network tab



NOTE:

Select models also include an option for configuring iLO settings.

Setup completion

After the storage server is physically set up and the basic configuration is established, additional setup steps must be completed. Depending on the deployment scenario of the storage server, these steps can vary.

Additional setup steps can include:

- Managing system storage
- Creating and managing users and groups
- Creating and managing file shares

Managing system storage

The storage server administrator uses Disk Management to manage volumes, and Shadow Copies to manage snapshots. See the following chapters for more detailed information on managing system storage:

- Chapter 3 discusses disk management procedures.
- Chapter 4 discusses snapshot (shadow copy) management procedures.
- Chapter 6 discusses folder and share management procedures.

Creating and managing users and groups

User and group information and permissions determine whether a user can access files. If the storage server is deployed into a workgroup environment, this user and group information is stored locally on the device. By contrast, if the storage server is deployed into a domain environment, user and group information is stored on the domain.

To enter local user and group information, see Chapter 5.

Creating and managing file shares

Files shares must be set up, granting and controlling file access to users and groups. See Chapter 6 for complete information on managing file shares.

UNIX specific information is discussed in the “[Services for NFS/UNIX](#)” chapter.



NOTE:

It is highly recommended to run Microsoft Windows Update to Identify, review, and install the latest, applicable, critical security update on the storage server. For recommendations, instructions, and documentation to help manage the software update, hotfix, and security patch processes on the storage server, see “Microsoft Software Updates on HP ProLiant Storage Servers” on the HP web site: <http://h18006.www1.hp.com/storage/storageservers.html>

Activating the iLO port using the license key

Select models include an iLO port. The Remote Desktop feature of the iLO port requires a license key. The key is included with the product inside the Country Kit. Refer to the iLO Advanced License Pack for activation instructions.

To access the iLO port, click **HP Utilities**, and then click **Remote Management**.

Setting up Ethernet NIC Teams (Optional)

Select models are equipped with the HP Network Teaming and Configuration utility. The utility allows administrators to configure and monitor Ethernet network interface controllers (NIC) teams in a Windows-based operating system. These teams provide options for increasing fault tolerance and throughput. See Appendix A for information on setting up NIC teams.

3 Disk and Volume Management

The process of creating storage elements and presenting them to the storage server OS is facilitated by the use of the WebUI.

This chapter presents information for storage servers with and without configuration storage.

Storage servers with configurable storage

The 300 series and 500 series storage servers ship pre-configured for the Operating System only. Additional storage configuration is needed. Depending on the type of storage server purchased, storage configuration may involve local storage configuration via the HP Array Configuration Utility or SAN management tools.

Select models of the 300 series and 500 series of storage servers also support clustering.

The primary web page for facilitating these storage servers is illustrated in [Figure 8](#). The diagram illustrates the process of creating arrays, volumes, and shadow copies. The steps on the left illustrate the logical steps used to manage disks, beginning with the array management at the top. The process follows the diagram from top to bottom and the selectable menu items from left to right on the page:

1. Create arrays and LUNS via the appropriate storage array management software.
2. Create disks via the WebUI.
3. Create volumes via the WebUI.

To manage disks and volumes via the WebUI, click **Disks**.




Figure 8 Disks menu—configurable storage models

The **Disks** tab contains the following task items:

Table 4 Disks tab options

Option	Task
Array Management	Open the Array Management page to access the ACU and links to other storage array management elements.
Disks	Manage logical disks. Observe disk capacity and status, scan for new disks, view detailed disk properties, and create new volumes.
Volumes	Manage disk space usage by enabling quotas, scheduling disk defragmentation, and performing detailed volume management using the Manage button.
Shadow Copies	Manage shadow copies of shared folders on the volume. Shadow copies are read-only copies of shared data that provide users via network shares with a way to view, and, if necessary, restore previous versions or deleted files.
Storage Guide	Provides a detailed list of the procedures required to configure and create disks and volumes on storage servers.

Storage configuration overview

On storage servers with configurable storage, physical disks can be arranged as RAID arrays for fault tolerance and enhanced performance, then segmented into logical disks of appropriate sizes for particular storage needs. These logical disks then become the volumes that appear as drives on the storage server.



NOTE:

This type of configuration may not apply to all supported storage components and serves only as an example.

The following steps are an example of a storage configuration using an HP Smart Array-based storage component.

Step 1: Create disk arrays

1. Click **Array Management** on the **Disk**s tab.
2. Click **Array Configuration Utility**, and then log in to the management page in another browser window.
The Array Configuration Utility starts.
3. Select the proper array controller in the left pane of the interface before beginning array configuration. Some storage servers are equipped with array controllers for both internal and external storage.

Consult the help available in ACU for details on creating arrays, if necessary.

Step 2: Create logical disks from the array space

From the ACU:

1. Select a previously created array.
2. Click **Create Logical Drive** in the right pane of the ACU.
3. Complete the Logical Drive Creation wizard to designate some or all of the array space as a logical disk.

Depending on how many physical disks are included in the array, several different types of logical disks are possible. Consult the ACU help for details on creating logical disk drives.

Step 3: Verify newly created logical disks

1. Click **Disk**s on the **Disk** tab.
2. Verify that disks matching the newly created sizes are displayed.
3. Click **Initialize Disk** to initialize the disk.



NOTE:

By default the disk is basic. Click Convert Disk to make the disk dynamic.

Step 4: Create a volume on the new logical disk

1. Click **Create New Volume**.

2. Enter the volume size.
3. Select a drive letter.
4. Enter a mount point, if desired.
5. Select to format the volume, if desired.
6. Click **OK**.
7. Select whether or not to quick format the volume.
8. Enter a volume label.
9. Enter the allocation unit size.
10. Click **OK**.

Array Configuration Utility (Smart Array-based storage only)

RAID arrays and LUNs on Smart Array-based storage servers are created and can be managed using the HP Array Configuration Utility.

Features of ACU:

- Graphical representation of drive array configurations with wizards that help optimize array configuration
- Online spare (hot spare) configuration
- Separate fault tolerance configurations on a logical drive (LUN) basis
- Easy capacity expansion of arrays
- Online RAID level and stripe size migration
- Manages OS and data drives

Each time the Array Configuration Utility is run, it analyzes the configuration of the Array Controllers installed in the system. From the Main page various options are available to change or reconfigure the controller(s). This document only covers a subset of the functions available in the ACU. For complete documentation on ACU, refer to the comprehensive online help found within the ACU tool.

Using the ACU to configure storage

To configure storage:

1. From the WebUI, click **Disks, Disks**.
2. Click **Array Management**.




Figure 9 Array Management page

3. Click **Array Configuration Utility**.



NOTE:

ACU is used to manage and configure array-based storage.

4. Log into the ACU. The default user name is administrator and the default password is administrator.




Figure 10 Systems Management Homepage

The Array Configuration Utility is displayed.

5. Select a controller in the list on the left side to begin configuration.
 - The first controller listed is for all drives in the server chassis, and may also connect to drives contained in an external storage enclosure attached to the server.
 - Additional controllers (if present) are used for all externally attached SCSI storage.



CAUTION:

On the first controller there are two logical drives pre-configured under Array A. These two logical drives are configured for the storage server operating system and should not be altered in any manner.

6. After the controller is selected there are three ways to configure the storage:
 - **Express Configuration**
Select **Express Configuration** to answer a few simple questions and allow the controller to be configured automatically. The **Express Configuration** is the easiest and fastest way to configure a controller, and provides the most reasonable configuration possible.
 - **Configuration Wizards**
Select **Configuration Wizards** to configure a controller through a series of wizards. Choosing **Configuration Wizards** is not the fastest or easiest way to configure a controller, but it does offer more control over the configuration and provides a more customized setup.
 - **Standard Configuration** (default)

Select **Standard Configuration** to quickly configure a controller. Choosing **Standard Configuration** is the fastest way to configure a controller but requires an intermediate to advanced level of knowledge concerning storage. The Standard Configuration path offers the least amount of help or step-by-step guides and does not provide a FAQ panel, assuming the user knows exactly what they would like to accomplish and are very familiar with the concepts required to complete the task.

The default method is the standard configuration method. The steps that follow are for creating an array using the standard configuration mode.

7. Click **Create Array**.
8. Select all of the drives to be included in the array, and then click **OK**.
9. Select the array that was just created, and then click **Create logical Drive** at the right.
10. Select the desired Fault Tolerance, Stripe Size, and Size of the logical disk, and then click **OK**.
The Fault tolerance level depends on the amount of disks selected when the array was created. A minimum of two disks is required for a RAID 0+1 configuration, three disks for a RAID 5 configuration, and four disks for a RAID 5 ADG configuration.
11. After all logical disks have been created, click **Save**.
12. Click **Exit ACU** to exit the ACU session.

ACU guidelines

- Do not modify Array A off of the Smart Array controller as it contains the storage server OS.
- Spanning more than 14 disks with a RAID 5 volume is not recommended.
- Designate spares for RAID sets to provide greater protection against failures.
- RAID sets cannot span controllers.
- A single array can contain multiple logical drives of varying RAID settings.
- Extending and Expanding Arrays and Logical Drives is supported.
- RAID migration is supported.

Managing disks on configurable storage servers

From the **Disks** tab of the WebUI, click **Disks**. The page displays the physical disks that are associated with the storage server and the volumes that are created on them. Multiple volumes can appear on multiple disks depending on whether the volumes are simple, spanned, or multi-volumes/partitions exist. The page also displays the type of disk (basic or dynamic).

The screenshot shows the 'Manage Disks' page of the StorageWorks interface. At the top, there's a navigation bar with links for Welcome, Status, Network, Disks, Users, Shares, Array Management, Maintenance, HP Utilities, Cluster, Help, and a question mark icon. Below the navigation is a sub-menu with links for Array Management, Disks, Volumes, Shadow Copies, Cluster Administrator, and Storage Guide. The main content area is titled 'Manage Disks' and contains a table with the following data:

Disk	Type	Volume(s)	Capacity	Status	Tasks
Disk 0	BASIC	(C:)	27 GB	Online	Rescan
Disk 1	BASIC	C:\WINDOWS\backup	7 GB	Online	Initialize Disk
Disk 2	BASIC		136 GB	Online	Disk Erase
Disk 3	BASIC		87 GB	Online	Convert Disk
Disk 4	BASIC		87 GB	Online	Create New Volume
Disk 5	BASIC		87 GB	Online	Properties
Disk 6	BASIC		87 GB	Online	Advanced Disk Management
Disk 7	BASIC	(L:)	87 GB	Online	
Disk 8	BASIC		1 GB	Online	

At the bottom left is a 'Back' button, and at the bottom right are 'Done' and 'Local intranet' buttons.

Figure 11 Manage Disks page—configurable storage server

Table 5 Manage Disks options

Option	Task
Rescan	Detects a new disk added to the system. By default, drives are dynamically recognized by the system. Occasionally a rescan of the hardware is required. The rescan is not synchronous and may require a browser refresh after the scan is initiated to display the new content.
Initialize Disk*	Initializes any empty disk to type basic.
Disk Erase*	Erases the selected disk.
Convert Disk*	Converts the selected disk from basic to dynamic, or dynamic to basic.
Create New Volume	Select to create a new volume.
Properties	Select to display the properties of the selected disk.
Advanced Disk Management	Select to open the Disk Management utility and perform advanced disk management tasks. Please see the online Disk Management help pages for complete documentation.

* These tasks cannot be completed on clustered resources.

Creating a new volume via the WebUI

To create a new volume via the WebUI:

1. Click the **Disks** tab, and then click **Disks**.
2. Select the Disk on which to create the new volume.
3. Click **Create New Volume**.




Figure 12 Creating a new volume, page 1

4. Enter the volume size.
5. Select a drive letter.
6. Enter a mount point, if desired.
7. Select to format the volume, if desired.
8. Click **OK**.




Figure 13 Creating a new volume, page 2

9. Select whether or not to quick format the volume.
10. Enter a volume label.
11. Enter the allocation unit size.
12. Click **OK**.

Advanced Disk Management

When **Advanced Disk Management** on the **Manage Disks** page is clicked, the Disk Management Utility is opened in a Remote Desktop session. The Disk Management tool is a system utility for managing hard disks and the volumes, or partitions, that they contain. The WebUI provides most of the functionality required for storage server disk management. However there are some instances where the Disk Manager is required. For example, to reassign a drive letter or mount point or to create software based RAID fault-tolerant disk systems.




Figure 14 Disk Management utility



NOTE:

When the Disk Management utility is accessed, the Remote Desktop connection assumes a dedicated mode and can only be used to manage disks and volumes on the server. Navigating to another page during an open session closes the session.



NOTE:

It may take a few moments for the Remote Desktop Connection session to log off when closing Disk Management.

Guidelines for managing disks

When managing disks and volumes:

- Read the online help found in the WebUI.
- Do not alter the Operating System Disk Labeled Local Disk C:.
- Do not alter the disk labeled “DON’T ERASE.”
- The use of software RAID-based dynamic volumes is not recommended; use the array controller instead, it is more efficient.
- The largest disk that Windows Storage Server 2003 can accommodate from a storage system is 2 TB.
- HP does not recommend spanning array controllers with dynamic volumes.
- Use meaningful volume labels with the intended drive letter embedded in the volume label, if possible. (For example, volume e: might be named “Disk E:.”) Volume labels often serve as the only means of identification.
- Record all volume labels and drive letters in case of system Quick Restore. (See “[Managing disks after quick restore](#)” later in this chapter).
- When managing basic disks, only the last partition on the disk can be extended unless the disk is changed to dynamic.
- Basic disks can be converted to dynamic without bringing the system offline or loss of data, but the volume is unavailable during the conversion.
- Basic disks can contain up to four primary partitions (or three primary partitions and one extended partition).
- Format drives with a 16 K allocation size for best support of shadow copies, performance, and defragmentation.
- NTFS formatted drives are recommended since they provide the greatest level of support for shadow copies, encryption, and compression.
- Only basic disks can be formatted as FAT or FAT32. Dynamic disks are not supported, nor can they be configured in a cluster.

Volumes page

On the Volumes page, administrators can select to manage volumes, schedule defragmentation, and set or manage quotas. The Volumes page displays all volumes that are formatted NTFS on the system. It does not display the volume type (for example simple or spanned) nor volumes that are FAT32 or FAT. To display these types of volumes, click **Manage**. All volumes are displayed.

See the **Managed Disks** page to view a list of disks, and the volumes assigned to them. The drive letters for volumes that encompass multiple disks appear on multiple rows on the display.




Figure 15 Volumes page

Table 6 Volumes page object/task selector

Option	Task
Manage...	Select to display the Manage Volumes page.
Schedule Defrag...	Select to schedule defragmentation for the selected volume.
Set Default Quota	Select to set quota limits to manage use of the volume. Settings on this page apply to new users and any users for whom user quota entries have not previously been set.
Set Quota Entries	Select to show a list of user quota entries. Then create a new quota entry, delete a quota entry, or view the properties of a quota entry.

Managing volumes

To manage volumes on the server:

1. From the WebUI, click the **Disks** tab.
2. Click **Volumes**.
3. In the Tasks list, click **Manage**.

The **Manage Volumes** page is displayed. The Manage Volumes page displays all volumes on the storage server regardless of their format (NTFS, FAT, or FAT32). Do not tamper with the "Don't Erase" or the Local C: volume. These are reserved volumes and must be maintained as they exist.

Volume Name	Type	Status	Capacity	Free Space	Tasks
Disk_M (M)	PARTITION	Healthy	1 GB	989 MB	Delete
DON'T ERASE (C:)	PARTITION	Healthy	7 GB	5540 MB	Extend
EDrive UnMounted	PARTITION	Healthy	87 GB	79839 MB	Rescan
FDrive (G:)	PARTITION	Healthy	87 GB	79838 MB	Advanced Volume Management
Local Disk (C:)	PARTITION	Healthy	27 GB	22565 MB	
Local Disk (I:)	PARTITION	Healthy	87 GB	88883 MB	
Quorum (Q:)	PARTITION	Healthy	136 GB	138853 MB	

Figure 16 Manage Volumes page

Table 7 Manage Volumes options

Option	Task
Delete*	Select to delete the selected volume. This is data destructive and there is no recovery other than from tape.
Extend	Opens a page to extend a partition based on a basic disk or to extend dynamic based volumes.
Rescan	Select to detect a volume or partition added to the system or to update the size of a volume that has undergone expansion. The rescan is not synchronous and may require a browser refresh after the scan is initiated to display the new content.
Advanced Volume Management	Select to open the Windows Disk Manager and perform advanced volume management tasks.

* This task cannot be completed on a clustered resource.

Dynamic growth

Dynamic growth is a feature of the storage server that provides for growth of a volume or partition to meet expanding storage requirements without the need to take volumes offline or incur downtime. Growth may occur in three forms:

- Extend unallocated space from original LUNS.
- Alter LUNs to contain additional storage.
- Add new LUNS to the system. The additional space is then extended through a variety of means, depending on which type of disk structure is in use.

Expanding a LUN

Expanding an existing LUN is accomplished using the storage array configuration software applicable to the storage array in use. In the case of the Smart Array controller, this is accomplished via the Array

Configuration Utility presented on the Disk page. LUN expansion may occur in Disk Arrays where space is available. If insufficient space is available, additional physical disks may be added to the array dynamically.




Figure 17 Expanding a LUN (Smart Array only)

To extend a LUN where space is available in the array (Smart Array only):

1. Click the **Disks** tab.
2. Click **Array Management**.
3. Click **Array Configuration Utility**, and then log in.
4. Select the appropriate array controller and the appropriate array that the logical drive is contained in.
5. Select the appropriate logical drive.
6. Click **Extend Size**.
7. Enter the total size of the logical drive in MB (not just the amount to be added), and then click **OK**.
8. Click **Save** to update the configuration.
9. Close the ACU.

To extend a LUN where space is not available in the array (Smart Array only):

1. Add an unassigned physical disk to the array using the ACU. If an unassigned physical disk is unavailable, add a new disk to the appropriate storage device, and then select **Refresh**.
2. To add an unassigned physical disk to the array use the following steps:
 - a. Select the appropriate array controller and the appropriate array that the logical drive is contained in.
 - b. Select **Expand Array**.

- c. Select the appropriate physical disk, and then click **OK**. The array is expanded.
3. Follow the instructions for extending a LUN.

Extending a partition on a basic disk

Partitions can be extended using either the WebUI extend function from the Managed Volumes page extend selection or by using the DiskPart command line utility. The Windows Disk Manager cannot extend basic disk partitions. To extend a partition using the WebUI follow the steps below:

1. From the WebUI, click the **Disk**s tab.
2. Click **Volumes**.
3. Click **Manage**.
4. Select the Volume to extend, and then click **Extend**.



NOTE:

If you receive a message that there is not enough disk space to extend the volume, it is possible to convert the disk to dynamic, provided that there are other dynamic disks with space available and that the storage server is not a node in a cluster. The volume can then be extended over a set of dynamic disks.

5. The page in [Figure 18](#) is displayed. Enter in the amount to extend the partition.

The screenshot shows a Microsoft Internet Explorer window titled "Extends the selected volume - Microsoft Internet Explorer". The address bar shows the URL as https://. The page header includes the HP StorageWorks logo and the Windows Storage Server 2003 logo. The main navigation menu has "Disks" selected. Below the menu, there's a sub-navigation bar with "Volumes" highlighted. The main content area is titled "Extend" and contains the following information:

Increase the size of a volume on a basic disk by adding space from the unallocated portion of the volume's underlying logical or physical disk.

Disk Name:	Disk 2
Volume Label:	New Volume
Volumes on Disk:	(E:)
Drive Letter:	(E:)
Disk Size:	13994 MB
Total disk space used:	5005 MB
Current Volume Size:	5005 MB
Disk Space Free:	8989 MB

Enter extend amount in Megabytes:

Volume Type: PARTITION
Volume Status: HEALTHY

At the bottom right of the dialog box are "OK" and "Cancel" buttons. At the very bottom of the browser window, there are "Done" and "Local intranet" buttons.

Figure 18 Extending a volume (basic disk)

6. Click **OK**.

Extending a volume on dynamic disks (non-clustered systems only)

The WebUI allows for the extension of volumes based on a dynamic disk or a set of dynamic disks. To extend a volume perform the following steps:

1. From the WebUI, click the **Disks** tab.
2. Click **Volumes**.
3. Click **Manage**.
4. Select the volume to extend, and then select **Extend**.




Figure 19 Extending a volume (dynamic disk)

5. Enter the amount to extend the volume. If no more space is available on the current dynamic disk, add an additional dynamic disk to the list of available disks and utilize space from it.
6. Click **OK**.

Extending using DiskPart

DiskPart may also be used to extend a partition or volume from the CMD prompt of the storage server operating system via Remote Desktop. Complete help is available from the Windows Storage Server 2003 desktop under **Start > Help and Support**. To use DiskPart follow the steps below:

Connect to the server through Remote Desktop, login, and then select the command prompt icon.

1. Enter Diskpart.
2. From the Diskpart prompt type the following commands:
 - Enter `list` to display all of the volumes.
 - Enter `select [name of volume]` (for example `select Volume 4`) to work against a particular volume or partition.

- Enter **Extend**. The volume is extended to the capacity of the underlying disk. To specify the amount to extend or to extend to another disk, enter: `extend [size=N] [disk=N]` Size is in MB.
- Enter **exit** to exit the utility.

Managing disks after quick restore

After a Quick Restore, drive letters may be assigned to the wrong volume. Windows Storage Server 2003 assigns drive letters after the restoration in the order of discovery. To help maintain drive letter information, placing the drive letter into the volume label is recommended. To change the drive letters to the appropriate ones, go into Disk Management and perform the following steps for each volume:

1. Right-click the volume that needs to be changed.
2. Select **Change Drive Letter and Paths**.
3. In the Change Drive Letter and Paths dialog box, select **Change**.
4. In the Change Drive Letter or Path dialog box, select the appropriate drive letter, and then click **OK**.
5. Click **Yes** to confirm the drive letter change.
6. Click **Yes** to continue. If the old drive letter needs to be reused, reboot the server after clicking **Yes**.

Storage servers with pre-configured storage

The primary WebUI page for facilitating disks and volume creation for storage servers with pre-configured storage is illustrated in [Figure 20](#). From this page the administrator can create and manage volumes via the WebUI.

To manage volumes via the WebUI, click **Disks**.




Figure 20 Disks tab—medium and small business class

The **Disks** tab contains the following task items for configuring the storage server:

Table 8 Disks tab options

Option	Task
Adaptec Storage Manager	Manage logical drives and view information about managed systems, controllers, disk groups, and so on.
Volumes	Manage disk space usage by enabling quotas, scheduling disk defragmentation, and performing detailed volume management using the Manage item.
Shadow Copies	Manage shadow copies of shared folders on the volume. Shadow copies are read-only copies of shared data that provide users with a way to view, and, if necessary, restore to previous versions of files.

Disk Management utility

When the **Advanced Volume Management** button on the Volumes page is selected, the Disk Management utility is opened after administrator login.

The Disk Management tool is a system utility for managing hard disks and the volumes, or partitions, that they contain. Disk Management is used to initialize disks, create volumes, format volumes with the FAT, FAT32, or NTFS file systems, and create fault-tolerant disk systems. Most disk-related tasks can be performed in Disk Management without restarting the system or interrupting users; most configuration changes take effect immediately. A complete online help facility is provided with the Disk Management Utility for assistance in using the product.




Figure 21 Disk Management utility



NOTE:

When the Disk Management utility is accessed, the Remote Desktop connection assumes a dedicated mode and can only be used to manage disks and volumes on the server. Navigating to another page during an open session closes the session.



NOTE:

It may take a few moments for the Remote Desktop Connection session to log off when closing Disk Management.

Disk Management guidelines

When managing disks and volumes:

- Read the online Disk Management help found in the utility.
- Do not alter the Operating System Disk labeled Primary OS C: and Secondary OS D:.
- Use meaningful volume labels with the intended drive letter embedded in the volume label, if possible. (For example, volume F: might be named “Disk F:.”) Volume labels often serve as the only means of identification.
- Record all volume labels and drive letters in case of system Quick Restore.
- When managing basic disks, only the last partition on the disk can be extended unless the disk is changed to dynamic.
- Basic disks can be converted to dynamic without bringing the system offline or loss of data, but the volume is unavailable during the conversion.
- Basic disks can contain up to four primary partitions (or three primary partitions and one extended partition).
- Format drives with a 16 K allocation size for best support of shadow copies, performance, and defragmentation.
- NTFS formatted drives are recommended since they provide the greatest level of support for shadow copies, encryption, and compression.
- Only basic disks can be formatted as FAT or FAT32.



NOTE:

The user guide contains Data Volume configuration information for the specific 100 series storage server.

Adaptec Storage Manager

Use the Adaptec Storage Manager to configure, administer, and monitor controllers that are installed locally or remotely in servers or storage enclosures. There is an extensive help system available in the application.



NOTE:

Storage on the 100 series storage servers is preconfigured at the factory. Access to Adaptec Storage Manager is for maintenance and monitoring.



NOTE:

Not all series 100 storage servers employ hardware RAID using the Adaptec Storage Manager.



NOTE:

The Adaptec Storage Manager version 2.12 uses Windows' user name and password for access and requires a login to access.




Figure 22 Adaptec Storage Manager

Volumes page

On the **Volumes** page, administrators can manage volumes, schedule defragmentation, and set or manage quotas. The Volumes page displays all volumes that are formatted NTFS on the system. It does not display the volume type (for example simple or spanned) nor volumes that are FAT32 or FAT. To display these types of volumes, click **Manage**.




Figure 23 Volumes tab

Table 9 Volumes page object/task selector

Option	Task
Advanced Volume Management	Select to display the Disk Management utility.
Schedule Defrag	Select to schedule defragmentation for the selected volume.
Set Default Quota	Select to set quota limits to manage use of the volume. Settings on this page apply to new users and any users for whom user quota entries have not previously been set.
Set Quota Entries	Select to show a list of user quota entries. Then create a new quota entry, delete a quota entry, or view the properties of a quota entry.

Scheduling defragmentation

The following information applies to all models of the HP ProLiant storage server.

Defragmentation is the process of analyzing local volumes and consolidating fragmented files and folders so that each occupies a single, contiguous space on the volume. This improves file system performance. Because defragmentation consolidates files and folders, it also consolidates the free space on a volume. This reduces the likelihood that new files will be fragmented.

Defragmentation for a volume can be scheduled to occur automatically at convenient times. Defragmentation can also be done once, or on a recurring basis.

To schedule defragmentation for a volume:

- From the WebUI, click the **Disks** tab.
- Click **Volumes**.
- Select the volume to schedule defragmentation.

4. In the Tasks list, click **Schedule Defrag**.
5. On the **Manage the defragmentation schedule for [VolumeName]** page, select the **Schedule defragmentation for this volume** check box.
6. Select the frequency: Once, Weekly, or Monthly.
7. Use the remaining controls to specify when defragmentation will occur. The available controls change according to the frequency that is selected.
8. Click **OK**.

To disable defragmentation for a volume:

1. From the WebUI, click the **Disks** tab.
2. Click **Volumes**.
3. Select the volume to disable defragmentation.
4. In the Tasks list, click **Schedule Defrag**.
5. On the **Manage the defragmentation schedule for [VolumeName]** page, clear the **Schedule defragmentation for this volume** check box.
6. Click **OK**.



NOTE:

Scheduling defragmentation to run no later than a specific time prevents the defragmentation process from running later than that time. If the defragmentation process is running when the time is reached, the process is stopped. This setting is useful to ensure that the defragmentation process ends before the demand for server access is likely to increase.

If defragmenting volumes on which shadow copies are enabled, use a cluster (or allocation unit) size of 16 KB or larger during the format. Otherwise defragmentation registers as a change by the Shadow Copy process. This increase in the number of changes forces Shadow Copy to delete snapshots as the limit for the cache file is reached.



CAUTION:

Allocation unit size cannot be altered without reformatting the drive. Data on a reformatted drive cannot be recovered.



NOTE:

NTFS compression is supported only if the cluster size is 4 KB or smaller.

Disk quotas

The following information applies to all models of the HP ProLiant Storage Server.

Disk quotas track and control disk space use in volumes.



NOTE:

To limit the size of a folder or share, see “[Directory quotas](#)” in Chapter 6.

Configure the volumes on the server to perform the following tasks:

- Prevent further disk space use and log an event when a user exceeds a specified disk space limit.
- Log an event when a user exceeds a specified disk space warning level.

When enabling disk quotas, it is possible to set both the disk quota limit and the disk quota warning level. The disk quota limit specifies the amount of disk space a user is allowed to use. The warning level specifies the point at which a user is nearing his or her quota limit. For example, a user's disk quota limit can be set to 50 megabytes (MB), and the disk quota warning level to 45 MB. In this case, the user can store no more than 50 MB on the volume. If the user stores more than 45 MB on the volume, the disk quota system logs a system event.

In addition, it is possible to specify that users can exceed their quota limit. Enabling quotas and not limiting disk space use is useful to still allow users access to a volume, but track disk space use on a per-user basis. It is also possible to specify whether or not to log an event when users exceed either their quota warning level or their quota limit.

When enabling disk quotas for a volume, volume usage is automatically tracked from that point forward, but existing volume users have no disk quotas applied to them. Apply disk quotas to existing volume users by adding new quota entries on the Quota Entries page.

Enabling quota management

When enabling disk quotas on a volume, every user's disk volume usage is monitored and treated differently, depending on the quota management settings for the specific user.

To enable quota management on a volume:

1. From the WebUI, click the **Disks** tab.
2. Click **Volumes**.
3. Select the volume to manage.
4. In the Tasks list, click **Set Default Quota**.
5. On the Default Quota for volume page, select **Use quota limits to manage use of the volume**.
6. If desired, select **Deny disk space to users exceeding quota limit** to enable that restriction.
7. Specify the default quota limit and warning level for new users on this volume.
8. Specify which quota events should be logged.
9. Click **OK**.



NOTE:

When enabling disk quotas on a volume, any users with write access to the volume who have not exceeded their quota limit can store data on the volume. The first time a user writes data to a quota-enabled volume, default values for disk space limit and warning level are automatically assigned by the quota system.

To disable quota management on a volume:

1. From the WebUI, click the **Disks** tab.
2. Click **Volumes**.
3. Select the volume to manage.
4. In the Tasks list, click **Set Default Quota**.
5. On the Default Quota for (volume) page, clear the check box to **Use quota limits to manage use of the volume**.
6. Click **OK**.

Setting user quota entries

The Set User Quotas page allows the administrator to set, delete, or change disk quotas for any user on the server.

To set or change quota entries on the server:

1. From the WebUI, click the **Disks** tab.
2. Click **Volumes**.
3. Select the volume to manage.
4. From the Tasks list, click **Set Quota Entries**.

The screenshot shows a user interface for managing user quotas. At the top, there's a navigation bar with links like Welcome, Status, Network, Disks, Users, Shares, Maintenance, HP Utilities, and Help. Below that is a sub-navigation bar for Adaptec Storage Manager, Volumes, and Shadow Copies. The main area is titled "Set User Quotas for Data (F:)" and contains a message: "You can set quota entries to manage volume use for individual users." A search bar at the top left allows filtering by "Login Name". Below it is a table with columns: Login Name, Status, Space Used, Usage Limit, Warning Level, and Tasks. One row is visible for "BUILTIN\Administrators" with values: No Limit, 0 Bytes, No Limit, No Limit. To the right of the table is a vertical "Tasks" column containing "New Quota Entry...", "Delete", and "Properties...". At the bottom left is a "Back" button.

Figure 24 Setting user quotas

To create a new user quota entry:

1. Click **New Quota Entry**.
2. Select a user.
3. Set the limit.
4. Set the warning level.
5. Click **OK**.

This screenshot shows the "Add New Quota Entry" dialog. It has tabs for "Disks" and "Volumes", with "Volumes" selected. The main area is titled "Add New Quota Entry". It asks to "Select a user from the list, or type a user login name in the text box." A list box on the left shows users: Administrator, ASPNET, Guest, IUSR_HPNAS, IWAM_HPNAS, and SUPPORT_388945a0. To the right is a text input field with placeholder text "Example: domain name\user". Below this are "Volume usage settings": a radio button for "Do not limit volume usage" (selected) and another for "Limit volume usage to:" followed by a numeric input field set to "No Limit" and a dropdown menu set to "KB". There's also a "Set warning level to:" section with similar controls.

Figure 25 Add new quota entry

To change a quota entry:

1. Select the quota to change.
2. Click **Properties**.
3. Change the limit.

4. Change the warning level.

5. Click **OK**.

To delete a quota entry:

1. Select the quota to change.

2. Click **Delete**.

DiskPart

The following information applies to all models of the HP ProLiant storage server.

DiskPart.exe is a text-mode command interpreter that enables the administrator to manage disks, partitions, or volumes.

When using the list commands, an asterisk (*) appears next to the object with focus. Select an object by its number or drive letter, such as disk 0, partition 1, volume 3, or volume C.

When selecting an object, the focus remains on that object until a different object is selected. For example, if the focus is set on disk 0 and volume 8 on disk 2 is selected, the focus shifts from disk 0 to disk 2, volume 8. Some commands automatically change the focus. For example, when creating a new partition, the focus automatically switches to the new partition.

Focus can only be given to a partition on the selected disk. When a partition has focus, the related volume (if any) also has focus. When a volume has focus, the related disk and partition also have focus if the volume maps to a single specific partition. If this is not the case, focus on the disk and partition is lost.

Table 10 Common DiskPart commands

Command	Description
add disk	Mirrors the simple volume with focus to the specified disk.
assign	Assigns a drive letter or mount point to the volume with focus.
convert basic	Converts an empty dynamic disk to a basic disk.
convert dynamic	Converts a basic disk into a dynamic disk. Any existing partitions on the disk become simple volumes.
create volume simple	Creates a simple volume. After creating the volume, the focus automatically shifts to the new volume.
exit	Exits the DiskPart command interpreter.
help	Displays a list of the available commands.
list disk	Displays a list of disks and information about them, such as their size, amount of available free space, whether the disk is a basic or dynamic disk, and whether the disk uses the master boot record (MBR) or GUID partition table. The disk marked with an asterisk (*) has focus.
list partition	Displays the partitions listed in the partition table of the current disk. On dynamic disks these partitions may not correspond to the dynamic volumes on the disk. This discrepancy occurs because dynamic disks contain entries in the partition table for the system volume or boot volume (if present on the disk). They also contain a partition that occupies the remainder of the disk in order to reserve the space for use by dynamic volumes.
list volume	Displays a list of basic and dynamic volumes on all disks.
rem	Provides a way to add comments to a script.
retain	Prepares an existing dynamic simple volume to be used as a boot or system volume.
select disk	Selects the specified disk and shifts the focus to it.



NOTE:

The Data Volume is configured by default as a RAID-5 volume across all four disks and is formatted as NFTS with a 16K allocation unit size.

For a complete list of DiskPart commands, go to the Windows Storage Server 2003 Desktop on the storage server via Remote Desktop and select **Start >Help and Support**, search on DiskPart.

Example of using DiskPart

The following example shows how to configure a volume on the storage server.

In the cmd window, type:

```
c:\>diskpart  
DISKPART>Rescan  
DISKPART>select disk 2
```

```
DISKPART>convert dynamic  
DISKPART>REM Create a simple volume  
DISKPART>create volume simple size=4000  
DISKPART> REM Assign drive letter F: to the volume  
DISKPART>assign letter=F  
DISKPART>list vol  
DISKPART>Exit
```

4 Shadow Copies

Overview



NOTE:

Select storage servers can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses using Shadow Copies in a non-clustered environment. See the Cluster Administration chapter of this guide for additional information regarding Shadow Copies in a cluster.

The Volume Shadow Copy Service provides an infrastructure for creating point-in-time snapshots (shadow copies) of volumes. Shadow Copy supports 64 shadow copies per volume.

A shadow copy contains previous versions of the files or folders contained on a volume at a specific point in time. While the Shadow Copy mechanism is managed at the server, previous versions of files and folders are only available over the network from clients and are seen on a per folder or file level and not as an entire volume.

The Shadow Copy feature uses data blocks. As changes are made to the file system, the Shadow Copy Service copies the original blocks to a special cache file, to maintain a consistent view of the file at a particular point in time. Because the snapshot only contains a subset of the original blocks, the cache file is typically smaller than the original volume. In the snapshot's original form, it takes up no space because blocks are not moved until an update to the disk occurs.

By using shadow copies, a storage server can maintain a set of previous versions of all files on the selected volumes. End users access the file or folder by using a separate client add-on program, which enables them to view the file in Windows Explorer. Accessing previous versions of files, or shadow copies, enables users to:

- Recover files that were accidentally deleted. Previous versions can be opened and copied to a safe location.
- Recover from accidentally overwriting a file. A previous version of that file can be accessed.
- Compare several versions of a file while working. Use previous versions to compare changes between two versions of a file.

Shadow copies cannot replace the current backup, archive, or business recovery system, but they can help to simplify restore procedures. Because a snapshot only contains a portion of the original data blocks, shadow copies can not protect against data loss due to media failures. However the strength of snapshots is the ability to instantly recover data from shadow copies, reducing the number of times needed to restore data from tape.

Shadow copy planning

Before setup is initiated on the server and the client interface is made available to end users, consider the following:

- From what volume will shadow copies be taken?
- How much disk space should be allocated for shadow copies?
- Will separate disks be used to store shadow copies?
- How frequently will shadow copies be made?

Identifying the volume

Shadow copies are taken for a complete volume, but not for a specific directory. Shadow copies work best when the server stores user files, such as documents, spreadsheets, presentations, graphics, or database files.



NOTE:

Shadow copies should not be used to provide access to previous versions of application or e-mail databases.

Shadow copies are designed for volumes that store user data such as home directories and My Documents folders that are redirected by using Group Policy or other shared folders in which users store data.

Shadow copies work with compressed or encrypted files and retain whatever permissions were set on the files when the shadow copies were taken. For example, if a user is denied permission to read a file, that user would not be able to restore a previous version of the file, or be able to read the file after it has been restored.

Although shadow copies are taken for an entire volume, users must use shared folders to access shadow copies. Administrators on the local server must also specify the \\servername\\sharename path to access shadow copies. If administrators or end users want to access a previous version of a file that does not reside in a shared folder, the administrator must first share the folder.



NOTE:

Shadow copies are available only on NTFS, not FAT or FAT32 volumes.

Files or folders that are recorded by using Shadow Copy appear static, even though the original data is changing.

Allocating disk space

When determining the amount of space to allocate for storing shadow copies, consider both the number and size of files that are being copied, as well as the frequency of changes between copies. For example, 100 files that only change monthly require less storage space than 10 files that change daily. If the frequency of changes to each file is greater than the amount of space allocated to storing shadow copies, no shadow copy is created.

Administrators should also consider user expectations of how many versions they will want to have available. End users might expect only a single shadow copy to be available, or they might expect three days or three weeks worth of shadow copies. The more shadow copies users expect, the more storage space administrators must allocate for storing them.

Setting the limit too low also affects backup programs that use shadow copy technology because these programs are also limited to using the amount of disk space specified by administrators.



NOTE:

Regardless of the volume space that is allocated for shadow copies, there is a maximum of 64 shadow copies for any volume. When the 65th shadow copy is taken, the oldest shadow copy is purged.

The minimum amount of storage space that can be specified is 350 megabytes (MB). The default storage size is 10 percent of the source volume (the volume being copied). If the shadow copies are stored on a separate volume, change the default to reflect the space available on the *storage* volume instead of the *source* volume. Remember that when the storage limit is reached, older versions of the shadow copies are deleted and cannot be restored.



CAUTION:

To change the storage volume, shadow copies must be deleted. The existing file change history that is kept on the original storage volume is lost. To avoid this problem, verify that the storage volume that is initially selected is large enough.

Converting basic storage disks to dynamic disks

When using a basic disk as a storage area for shadow copies and converting the disk into a dynamic disk, it is important to take the following precaution to avoid data loss:

- If the disk is a non-boot volume and is a different volume from where the original files reside, first dismount and take offline the volume containing the original files before converting the disk containing shadow copies to a dynamic disk.
- The volume containing the original files must be brought back online within 20 minutes, otherwise, the data stored in the existing shadow copies is lost.
- If the shadow copies are located on a boot volume, the disk to can be converted to dynamic without losing shadow copies.



NOTE:

Use the mountvol command with the /p option to dismount the volume and take it offline. Mount the volume and bring it online using the mountvol command or the Disk Management snap-in.

Identifying the storage area

To store the shadow copies of another volume on the same file server, a volume can be dedicated on separate disks. For example, if user files are stored on *H:*, another volume such as *S:* can be used to store the shadow copies. Using a separate volume on separate disks provides better performance and is recommended for heavily used storage server.

If a separate volume will be used for the storage area (where shadow copies are stored), the maximum size should be changed to **No Limit** to reflect the space available on the storage area volume instead of the source volume (where the user files are stored).

Disk space for shadow copies can be allocated on either the same volume as the source files or a different volume. There is a trade-off between ease of use and maintenance versus performance and reliability that the system administrator must consider.

By keeping the shadow copy on the same volume there is a potential gain in ease of setup and maintenance; however there may be a reduction in performance and reliability.



CAUTION:

If shadow copies are stored on the same volume as the user files, note that a burst of disk input/output (I/O) can cause all shadow copies to be deleted. If the sudden deletion of shadow copies is unacceptable to administrators or end users, it is best to use a separate volume on separate disks to store shadow copies.

Determining creation frequency

The more frequently shadow copies are created, the more likely that end users will get the version that they want. However, with a maximum of 64 shadow copies per volume, there is a trade-off between the frequency of making shadow copies and the amount of time that the earlier files will be available.

By default, the storage server creates shadow copies at 0700 and 1200, Monday through Friday. However, these settings are easily modified by the administrator so that the shadow copy schedule can better accommodate end user needs. To modify these schedules see “[Scheduling shadow copies](#)” later in this chapter.

Shadow copies and drive defragmentation

When running Disk Defragmenter on a volume with shadow copies activated, all or some of the shadow copies may be lost, starting with the oldest shadow copies.

If defragmenting volumes on which shadow copies are enabled, use a cluster (or allocation unit) size of 16 KB or larger. Using this allocation unit size reduces the number of copy outs occurring on the snapshot. Otherwise the number of changes caused by the defragmentation process can cause shadow copies to be deleted faster than expected. Note, however, that NTFS compression is supported only if the cluster size is 4 KB or smaller.



NOTE:

To check the cluster size of a volume, use the `fsutil fsinfo ntfsinfo` command. To change the cluster size on a volume that contains data, back up the data on the volume, reformat it using the new cluster size, and then restore the data.

Mounted drives

A mounted drive is a local volume attached to an empty folder (called a mount point) on an NTFS volume. When enabling shadow copies on a volume that contains mounted drives, the mounted drives are not included when shadow copies are taken. In addition, if a mounted drive is shared and shadow copies are enabled on it, users cannot access the shadow copies if they traverse from the host volume (where the mount point is stored) to the mounted drive.

For example, assume there is a folder *F:\data\users*, and the *Users* folder is a mount point for *G:*. If shadow copies are enabled on both *F:* and *G:*, *F:\data* is shared as *\server1\data*, and *G:\data\users* is shared as *\server1\users*. In this example, users can access previous versions of *\server1\data* and *\server1\users* but not *\server1\data\users*.

Managing shadow copies

From the **WebUI Welcome** screen, click **Disks**, and then **Shadow Copies** to display the **Shadow Copies** page.

Figure 26 Shadow Copies page

Table 11 Shadow Copies fields

Field	Description
Volume	Lists all volumes of the server on which the Shadow Copies service can be used. Only NTFS file system data volumes that are physically located on the server can support shadow copies. To manage shadow copies on a volume, select the check box next to the volume name, and then choose a task from the Tasks list.
Copies	Lists the number of shadow copies on the volume.
Used Space	Lists the total disk space that is used by the shadow copies on the volume.
Shares	Lists the number of shared folders that reside on the volume. This information can help determine whether to enable shadow copies on a volume. A greater number of shared folders on a volume increases the likelihood that users might need access to previous versions of their data.
Next Run Time	If the Shadow Copies service is enabled on the volume, this column lists the time and date the next shadow copy will be created. Otherwise, it displays Disabled.

Table 12 Shadow Copies tasks

Task	Click to...
Enable	enable Shadow Copies on the selected volume.
Disable	disable Shadow Copies on the selected volume.
New Shadow Copy	immediately create a new shadow copy on the selected volume.
View Shadow Copies	view a list of shadow copies on the selected volume.
Set Schedule	set the time and frequency of shadow copies.
Properties...	view the shadow copy properties of the selected volume, including location and size of the cache file.
Manage Shares	go to the Shared Folders page.

The shadow copy cache file

The default shadow copy settings allocate 10 percent of the source volume being copied (with a minimum of 350 MB), and store the shadow copies on the same volume as the original volume. (See [Figure 27](#)). The cache file is located in a hidden protected directory entitled "System Volume Information" off of the root of each volume for which Shadow Copy is enabled.




Figure 27 Shadow copies stored on source volume

The cache file location can be altered to reside on a dedicated volume separate from the volumes containing files shares. (See [Figure 28](#)).




Figure 28 Shadow copies stored on separate volume

The main advantage to storing shadow copies on a separate volume is ease of management and performance. Shadow copies on a source volume must be continually monitored and can consume space designated for file sharing. Setting the limit too high takes up valuable storage space. Setting the limit

too low can cause shadow copies to be purged too soon, or not created at all. By storing shadow copies on a separate volume space, limits can generally be set higher, or set to No Limit. See the **Properties** tab of the **Shadow Copy** page for a volume to alter the cache file location.



CAUTION:

If the data on the separate volume L: is lost, the shadow copies cannot be recovered.

Enabling and creating shadow copies

Enable the Shadow Copies service for a volume or create a shadow copy directly from the **Manage Shadow Copies** page.

Enabling shadow copies on a volume automatically results in several actions:

- Creates a shadow copy of the selected volume
- Sets the maximum storage space for the shadow copies
- Schedules shadow copies to be made at 7 A.M. and 12 noon on weekdays



NOTE:

Creating a shadow copy only makes one copy of the volume; it does not create a schedule.

To enable shadow copies on a volume:

1. From the WebUI, click **Disks**.
2. Click the **Shadow Copies** tab.
3. On the **Manage Shadow Copies** page, select one or more volumes to enable the Shadow Copies service on.



NOTE:

After the first shadow copy is created, it cannot be relocated. Relocate the cache file by altering the cache file location under Properties prior to enabling shadow copy. See “[Viewing shadow copy properties](#)” in this chapter.

4. Click **Enable**.

To create a shadow copy on a volume:

1. From the WebUI, click **Disks**.
2. Click the **Shadow Copies** tab.
3. On the **Manage Shadow Copies** page, select one or more volumes to create the shadow copies on.
4. Click **New Shadow Copy**.

Viewing a list of shadow copies

To view a list of shadow copies on a volume:

1. From the WebUI, click **Disks**.
2. Click the **Shadow Copies** tab.
3. On the **Manage Shadow Copies** page, select the volume to view.
4. On the Tasks list, click **View Shadow Copies**.

All shadow copies are listed, sorted by the date and time they were created.



NOTE:

It is also possible to create new shadow copies or delete shadow copies from this page.

Set schedules

Shadow Copy schedules control how frequently shadow copies of a volume are made. There are a number of factors that can help determine the most effective shadow copy schedule for an organization. These include the work habits and locations of the users. For example, if users do not all live in the same time zone, or they work on different schedules, it is possible to adjust the daily shadow-copy schedule to allow for these differences.

Do not schedule shadow copies more frequently than once per hour.

Scheduling shadow copies

When the Shadow Copies service is enabled on a volume, it automatically schedules shadow copies to be made each weekday at 7 A.M. and 12 noon.

To add or change a shadow copy schedule for a volume:

1. From the WebUI, click **Disks**.
2. Click **Shadow Copies**.
3. Select the volume.
4. In the Tasks list, click **Set Schedule**.
5. On the **Shadow Copy Schedules** page, click **New**.
6. Select a frequency: Once, Daily, Weekly, or Monthly.
7. Use the remaining controls to specify the recurrence pattern and the starting date and time. The available controls change according to the frequency selected.
8. Click **OK**.

Deleting a shadow copy schedule

To delete a shadow copy schedule on a volume:

1. From the WebUI, click **Disks**.
2. Click the **Shadow Copies** tab.
3. Select the volume on which to delete a shadow copy schedule.
4. In the Tasks list, click **Set Schedule**.
5. On the **Manage Shadow Copy Schedules** page, select the schedule to be deleted, and then click **Delete**.
6. Click **OK** to confirm the deletion or **Cancel** to retain the copy.



NOTE:

When deleting a shadow copy schedule, that action has no effect on existing shadow copies. To remove schedules and all shadow copies in one action, from the Manage Shadow Copies page, click **Disable** on the Tasks list.

Viewing shadow copy properties

To view shadow copy properties on a volume:

1. From the WebUI, click **Disks**.
2. Click the **Shadow Copies** tab.
3. On the **Manage Shadow Copies** page, select the volume on which to view shadow copy properties.
4. In the Tasks list, click **Properties**.

The **Shadow Copy Properties** page, as shown in [Figure 29](#), lists the number of copies, the date and time the most recent shadow copy was made, and the maximum size setting.

Change the maximum size limit for all shadow copies, or choose **No limit**.

For volumes where shadow copies do not exist currently, it is possible to change the location of the cache file. (See “[The shadow copy cache file](#)” earlier in this chapter). The list of available disks and the space available on each is presented at the bottom of the page. Managing the cache files on a separate disk is recommended.



NOTE:

If shadow copies have already been enabled, the cache file location is grayed out. To change this location after shadow copies have been enabled, all shadow copies must be deleted and cannot be recovered. Remember enabling Shadow Copies creates a shadow copy by default.

5. Click **OK** to save changes, or click **Cancel** to discard changes.




Figure 29 Shadow Copy Properties page



CAUTION:

Use caution when reducing the size limit for all shadow copies. When the size is set to less than the total size currently used for all shadow copies, enough shadow copies are deleted to reduce the total size to the new limit. A shadow copy cannot be recovered after it has been deleted.

Disabling shadow copies

When shadow copies are disabled on a volume, all existing shadow copies on the volume are deleted as well as the schedule for making new shadow copies.

To disable shadow copies on a volume:

1. From the WebUI, click **Disks**.
2. Click the **Shadow Copies** tab.
3. On the **Manage Shadow Copies** page, select one or more volumes on which to disable shadow copies.
4. In the Tasks list, click **Disable**.

The **Disable Shadow Copies** page identifies the volume for which shadow copies will be disabled.

5. Click **OK** to delete all existing shadow copies and settings for the volume.



CAUTION:

When the Shadow Copies service is disabled, all shadow copies on the selected volumes are deleted. Once deleted, shadow copies cannot be restored.

Managing shadow copies from the storage server desktop

As an alternative to managing Shadow Copies via the WebUI, the storage server desktop can be accessed via Remote Desktop.

To access Shadow Copies from the storage server desktop:

1. From the WebUI, click **Maintenance, Remote Desktop**.
2. Click **My Computer**.
3. Select the volume.
4. Right-click the volume name and select **Properties**.
5. Click the **Shadow Copies** tab.

The user interface provides the same functionality found in the WebUI but in Win32 form. See [Figure 30](#).




Figure 30 Accessing Shadow Copies from My Computer

Shadow Copies for Shared Folders

Shadow Copies are accessed over the network by supported clients and protocols. There are two sets of supported protocols, SMB and NFS. All other protocols are not supported; this would include HTTP, FTP, AppleTalk, and NetWare Shares. For SMB support a client side application denoted as Shadow Copies for Shared Folders is required. The client side application is currently only available for Windows XP and Windows 2000 SP3+. The application is included on the storage server device from the following directory:

C:\hpnas\Components\ShadowCopyClient\XP and 2000-SP3+

No additional software is required to enable UNIX users to independently retrieve previous versions of files stored on NFS shares.



NOTE:

Shadow Copies for Shared Folders supports retrieval only of shadow copies of network shares. It does not support retrieval of shadow copies of local folders.



NOTE:

Shadow Copies for Shared Folders clients are not available for HTTP, FTP, AppleTalk, or NetWare shares. Consequently, users of these protocols cannot use Shadow Copies for Shared Folders to independently retrieve previous versions of their files. However, administrators can take advantage of Shadow Copies for Shared Folders to restore files on behalf of these users.

SMB shadow copies

Windows users can independently access previous versions of files stored on SMB shares via the Shadow Copies for Shared Folders client. After the Shadow Copies for Shared Folders client is installed on the user's computer, the user can access shadow copies for a share by right-clicking on the share to open its Properties window, clicking the **Previous Versions** tab, and then selecting the desired shadow copy. Users can view, copy, and restore all available shadow copies.

Shadow Copies for Shared Folders preserves the permissions set in the access control list (ACL) of the original folders and files. Consequently, users can only access shadow copies for shares to which they have access. In other words, if a user does not have access to a share, he also does not have access to the share's shadow copies.

The Shadow Copies for Shared Folders client pack installs a **Previous Versions** tab in the **Properties** window of files and folders on network shares.

Users access shadow copies with Windows Explorer by selecting **View**, **Copy**, or **Restore**, from the **Previous Versions** tab. (See [Figure 31](#)). Both individual files and folders can be restored.




Figure 31 Client GUI

When users view a network folder hosted on the storage server for which shadow copies are enabled, old versions (prior to the snapshot) of a file or directory are available. Viewing the properties of the file or folder presents users with the folder or file history—a list of read-only, point-in-time copies of the file or folder contents that users can then open and explore like any other file or folder. Users can view files in the folder history, copy files from the folder history, and so on.

NFS shadow copies

UNIX users can independently access previous versions of files stored on NFS shares via the NFS client; no additional software is required. Server for NFS exposes each of a share's available shadow copies as a pseudo-subdirectory of the share. Each of these pseudo-subdirectories is displayed in exactly the same way as a regular subdirectory is displayed.

The name of each pseudo-subdirectory reflects the creation time of the shadow copy, using the format `@GMT-YYYY.MM.DD-HH:MM:SS`. To prevent common tools from needlessly enumerating the pseudo-subdirectories, the name of each pseudo-subdirectory begins with the dot character, thus rendering it hidden.

The following example shows an NFS share named “NFSShare” with three shadow copies, taken on April 27, 28, and 29 of 2003 at 4 a.m.

NFSShare

```
.@GMT-2003.04.27-04:00:00  
.@GMT-2003.04.28-04:00:00  
.@GMT-2003.04.29-04:00:00
```

Access to NFS shadow copy pseudo-subdirectories is governed by normal access-control mechanisms using the permissions stored in the file system. Users can access only those shadow copies to which they have read access at the time the shadow copy is taken. To prevent users from modifying shadow copies, all pseudo-subdirectories are marked read-only, regardless of the user's ownership or access rights, or the permissions set on the original files.

Server for NFS periodically polls the system for the arrival or removal of shadow copies and updates the root directory view accordingly. Clients then capture the updated view the next time they issue a directory read on the root of the share.

Recovery of files or folders

There are three common situations that may require recovery of files or folders:

- Accidental file deletion, the most common situation.
- Accidental file replacement, which may occur if a user selects Save instead of Save As.
- File corruption.

It is possible to recover from all of these scenarios by accessing shadow copies. There are separate steps for accessing a file compared to accessing a folder.

Recovering a deleted file or folder

To recover a deleted file or folder within a folder:

1. Navigate to the folder where the deleted file was stored.
2. Position the cursor over a blank space in the folder. If the cursor hovers over a file, that file is selected.
3. Right-click, select **Properties** from the bottom of the menu, and then click the **Previous Versions** tab.
4. Select the version of the folder that contains the file before it was deleted, and then click **View**.
5. View the folder and select the file or folder to recover. The view may be navigated multiple folders deep.
6. Click **Restore** to restore the file or folder to its original location. Click **Copy...** to allow the placement of the file or folder to a new location.




Figure 32 Recovering a deleted file or folder

Recovering an overwritten or corrupted file

Recovering an overwritten or corrupted file is easier than recovering a deleted file because the file itself can be right-clicked instead of the folder. To recover an overwritten or corrupted file:

1. Right-click the overwritten or corrupted file, and then click **Properties**.
2. Click **Previous Versions**.
3. To view the old version, click **View**. To copy the old version to another location, click **Copy...** to replace the current version with the older version, click **Restore**.

Recovering a folder

To recover a folder use the following procedure:

1. Position the cursor so that it is over a blank space in the folder to be recovered. If the cursor hovers over a file, that file is selected.
2. Right-click, select **Properties** from the bottom of the menu, and then click the **Previous Versions** tab.
3. Click either **Copy...** or **Restore**.

Clicking **Restore** enables the user to recover everything in that folder as well as all subfolders. Clicking **Restore** does not delete any files.

Backup and shadow copies

Shadow copies are only available on the network via the client application and only at a file or folder level as opposed to the entire volume. Hence the standard backup associated with a volume backup will not work to back up the previous versions of the file system. To answer this particular issue, shadow copies are available for back up in two situations. If the backup software in question supports the use of shadow copies and can communicate with underlying block device, it is supported and the previous version of the file system will be listed in the backup application as a complete file system snapshot. Lastly, if the built-in backup application NTbackup is utilized, the backup software forces a snapshot and then uses the snapshot as the means for back up. The user is unaware of this activity and it is not self evident although it does address the issue of open files.

5 User and Group Management

Overview

There are two system environments for users and groups: workgroup and domain. Because users and groups in a domain environment are managed through standard Windows or Active Directory domain administration methods, this document discusses only local users and groups, which are stored and managed on the storage server. For information on managing users and groups on a domain, refer to the domain documentation available on the Microsoft web site.

Domain compared to workgroup environments

When a storage server is deployed into a workgroup environment, all user and group account access permissions to file resources are stored locally on the server.

By contrast, when a storage server is deployed into a domain environment it uses the account database from the domain controller, with user and group accounts stored outside the server. The server integrates with the domain controller infrastructure.



NOTE:

The storage server cannot act as a domain controller for other servers on the network. If user and group account information is stored locally, those accounts may be used only to authenticate logons to the storage server, resulting in a workgroup configuration.

Additional information about planning for domain environments can be found at:

<http://www.microsoft.com/windowsserver2003/technologies/directory/activedirectory/default.mspx>

User and group name planning

Effective user and group management depends upon how well the user and group names are organized. Administrators typically create a small number of groups on the network and then assign users to the appropriate group or groups. File system and share permissions can then be applied at the group level, rather than at the user level. If the number of groups is small, assigning the appropriate permissions to the selected group, or groups, is more efficient than assigning permissions to each user.

Although each organization has specific conventions, following general guidelines makes administration simpler and more efficient. Because CIFS/SMB is dependent on users and groups to grant appropriate access levels to file shares, CIFS/SMB administration benefits from a consistent user and group administration strategy.

Managing user names

Usernames should reflect a logical relationship between the username and the person who uses the account. It is important that rules are established to ensure that usernames are:

- Systematic.
- Easy to follow and implement.
- Easy to remember.

Using a combination of the user's first name, middle initial, and last name results in systematic usernames for every member of a particular organization. For example, first initial followed by last name (jdoe for John Doe).

Guidelines must be in place for instances when two users have the same initials or name. For example, a number can be added to the end of the username (jdoe1 and jdoe2).

Other conventions can be applied. Just ensure that conventions are both systematic and consistent.

Managing group names

Group management follows many of the same principles as user management.

Group naming conventions should be systematic and easy to understand. Make the group name convey some logical information about the function or purpose of the group. [Table 13](#) provides examples of group names.

Table 13 Group name examples

Group Name	Description
Administrators	All designated administrators on the server
Users	All standard server users
Power users	All standard server users requiring advanced access levels

Using tags is a helpful convention that indicates the specific access that a particular user has to a network resource. For example, if there is a data share on the device, the network administrator can create a "Data Users ROnly" group and a "Data Users RWWrite" group to contain users that have read only or read write access on the share, respectively.

Workgroup user and group management



NOTE:

In a clustered environment, users and groups should not be managed locally.

In a workgroup environment, local users and groups are managed through the WebUI of the storage server.

Managing local users

In the WebUI, click **Users, Local Users** to display the **Local Users on Server** page. All workgroup user administration tasks are performed here.

The screenshot shows the 'Local Users' page with the following interface elements:

- Header:** 'Users' tab is selected.
- Navigation:** 'Local Users' and 'Local Groups' tabs.
- Title Bar:** 'Local Users on Server'.
- Message:** 'Select a user, then choose a task. To create a new user, choose New...'. A search bar with 'Name' dropdown and 'Go' button is present.
- Table:** Displays a list of local users with columns: 'Name' (checkbox), 'Full Name', 'Account is disabled', and 'Tasks' (button).
- Tasks Column:** Contains buttons for 'New...', 'Delete', 'Set a Password...', and 'Properties...'.
- Buttons:** 'Back' button at the bottom left.

Name	Full Name	Account is disabled	Tasks
<input type="checkbox"/> Administrator		No	
<input type="checkbox"/> Guest		Yes	
<input type="checkbox"/> IUSR_NAS	Internet Guest Account	No	
<input type="checkbox"/> IWAM_NAS	Launch IIS Process Account	No	
<input type="checkbox"/> sfuuser	sfuuser	No	
<input type="checkbox"/> TslInternetUser	TslInternetUser	No	

Figure 33 Local Users page

When the **Local Users** page is initially displayed, only the **New** option is available. After an existing user is selected, the additional actions are displayed. Existing user records can be retrieved in one of two ways:

- Enter the User Name or Full Name in the Search fields to retrieve a specific user record. To redisplay the complete user list, space out the Search field.
- Select the user in the list of displayed users in the page. The sort order of the display is controlled by clicking the Name field heading. The names are displayed in alphanumeric order or reverse alphanumeric order.

Adding a new user

To add a user:

1. On the **Local Users** page, click **New**.

The screenshot shows the 'Create New User' interface. At the top, there's a navigation bar with 'Local Users' and 'Local Groups' tabs, and a 'Users' tab which is active. Below this is a blue header bar with the title 'Create New User'. The main content area is titled 'Create New User' and contains several input fields: 'User name', 'Full name', 'Description', 'Password', 'Confirm password', and 'Home Directory'. Under 'Home Directory', there's a checkbox labeled 'Path' followed by a text input field showing a path. Below this are three checkboxes: 'Disable this user account', 'Password never expires', and another one that is currently checked.

Figure 34 Create New User page

2. Enter the user information, and then click **OK**.

Deleting a user

To delete a user:

1. On the **Local Users** page, select the user to delete, and then click **Delete**.

The **Delete User** dialog box is displayed, including a warning note about deleting users.

2. To delete the user, click **OK**.

Modifying a user password

To modify a user password:

1. On the **Local Users** page, select the user whose password needs to be changed.
2. Click **Set a Password**.
3. Enter the password, and then click **OK**.

Modifying user properties

To modify other user properties:

1. On the **Local Users** page, select the user whose record needs to be modified.
2. Click **Properties**.




Figure 35 User Properties page

3. Complete the changes, and then click **OK**.

Managing local groups

In the WebUI, click **Users, Local Groups** to display the **Local Groups on Server** page.

Name	Description
Administrators	Administrators have complete and unrestricted access to all resources on the system.
Backup Operators	Backup Operators can override security restrictions to back up files and folders.
Guests	Guests have the same access as members of the User group.
PasswordPropDeny	Users whose passwords should not be synchronized.
Power Users	Power Users possess most administrative powers without being Administrators.
Replicator	Supports file replication in a domain.
TelnetClients	TelnetClients users can access Windows NT/2000 systems.
Users	Users are prevented from making accidental or intentional changes to their accounts.

Figure 36 Local Groups page

Adding a new group

To add a group:

1. On the **Local Groups** page, click **New**.




Figure 37 Create New Group page, General tab

2. Enter the group name and description.
3. To indicate the user members of this group, click **Members**. See “[Modifying group properties](#)” for additional information.
4. After all group information is entered, click **OK**.

Deleting a group

To delete a group:

1. On the **Local Groups** page, select the group to delete, and then click **Delete**.
2. The **Delete Group** page is displayed.
3. Verify that this is the intended group, and then click **OK**.

Modifying group properties

To modify other group properties:

1. On the **Local Groups** page, select the desired group, and then click **Properties**.




Figure 38 Group Properties page, General tab

2. Enter the desired changes in each of the tabs, and then click **OK**.

General tab

Use the **General** tab to change basic group information, including:

- Group name
- Description

Members tab

Within the **Members** tab, users are added and removed from a group.

Current members of the group appear in the **Members** box. All users are listed in the **Add user or group** box.

- To add an existing local user to a group:
 1. Select the desired user from the **Add user or group** box.
 2. Click **Add**.
 3. Click **OK**.
- To remove an existing local user from a group:
 1. Select the desired user from the **Members** box.
 2. Click **Remove**.
 3. Click **OK**.
- To add a domain user or group to this group:
 1. Enter the user or group name to include in the indicated format (domain/username).
 2. Select **Add**.
 3. Enter a domain/username and password.
 4. Click **OK**.



NOTE:

To add domain users and groups to a local group, the storage server must be a member of the domain.




Figure 39 Group Properties page, Members tab

6 Folder, Printer, and Share Management

The HP ProLiant Storage Server supports several file sharing protocols, including DFS, NFS, FTP, HTTP, and Microsoft SMB. This chapter discusses overview information as well as procedural instructions for the setup and management of the file shares for the supported protocols. Security at the file level and at the share level are also discussed.

Abbreviated information on creating NFS file shares is included in this chapter; for detailed information on setting up and managing NFS file shares, see the “[Services for NFS/UNIX](#)” chapter.

NCP shares must be set up and managed through the Management Console user interface. For information on managing NCP file shares, see the “[NetWare File System Management](#)” chapter.

More information about Windows file system security is available on the Microsoft web site:

www.microsoft.com/



NOTE:

Select servers can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses share setup for a non-clustered deployment. For information on managing file shares and printers in a cluster, see the “[Cluster Administration](#)” chapter.

Folder management

Volumes and folders on any system are used to organize data. Regardless of system size, systematic structuring and naming conventions of volumes and folders eases the administrative burden. Moving from volumes to folders to shares increases the level of granularity of the types of data stored in the unit and the level of security access allowed.

Although a variety of methods can be used to create and manage file folders on the storage server, this document discusses using the web-based user interface (WebUI.) For additional information, use the WebUI online help.

Managing system volumes and file folders includes the following tasks:

Navigating to a specific volume or folder

When you work with volumes and folders, the first task is to gain access to the desired volume or folder.

The steps are the same, whether navigating to a volume or a folder:

1. From the WebUI, click **Shares**, and then **Folders**.

Shares																																																						
Folders Shares Sharing Protocols Directory Quota Storage Reports File Screening																																																						
 Folders																																																						
You can manage the folders and shares for each volume.																																																						
<table><tr><td>Search: <input type="text" value="Volume Name"/></td><td><input type="button" value="Go"/></td><td><input type="button" value=""/></td><td><input type="button" value=""/></td><td></td></tr><tr><th><input type="checkbox"/> Volume Name</th><th>Total Size</th><th>Free Space</th><th>Share Type</th><th>Tasks</th></tr><tr><td><input type="checkbox"/> Data (F)</td><td>671 GB</td><td>671 GB</td><td></td><td><input type="button" value="Manage Folders"/></td></tr><tr><td><input type="checkbox"/> Primary OS (C)</td><td>8.99 GB</td><td>5.74 GB</td><td></td><td><input type="button" value="Share Volume..."/></td></tr><tr><td><input type="checkbox"/> Secondary OS (D)</td><td>8.99 GB</td><td>5.82 GB</td><td></td><td><input type="button" value="Manage Shares..."/></td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr></table>					Search: <input type="text" value="Volume Name"/>	<input type="button" value="Go"/>	<input type="button" value=""/>	<input type="button" value=""/>		<input type="checkbox"/> Volume Name	Total Size	Free Space	Share Type	Tasks	<input type="checkbox"/> Data (F)	671 GB	671 GB		<input type="button" value="Manage Folders"/>	<input type="checkbox"/> Primary OS (C)	8.99 GB	5.74 GB		<input type="button" value="Share Volume..."/>	<input type="checkbox"/> Secondary OS (D)	8.99 GB	5.82 GB		<input type="button" value="Manage Shares..."/>																									
Search: <input type="text" value="Volume Name"/>	<input type="button" value="Go"/>	<input type="button" value=""/>	<input type="button" value=""/>																																																			
<input type="checkbox"/> Volume Name	Total Size	Free Space	Share Type	Tasks																																																		
<input type="checkbox"/> Data (F)	671 GB	671 GB		<input type="button" value="Manage Folders"/>																																																		
<input type="checkbox"/> Primary OS (C)	8.99 GB	5.74 GB		<input type="button" value="Share Volume..."/>																																																		
<input type="checkbox"/> Secondary OS (D)	8.99 GB	5.82 GB		<input type="button" value="Manage Shares..."/>																																																		

Figure 40 Volumes page

2. Select the appropriate volume, and then click **Manage Folders** to display a list of all of the folders within that volume.
3. To navigate to a subfolder, select the folder in which the subfolder resides, and then click **Open**. Repeat this searching and opening process until the desired folder is opened.

Figure 41 Folders page

Creating a new folder

To create a new folder:

1. Click the **Shares** tab, **Folders**, **Manage Folders**, and then click **New**.
2. On the **General** tab, enter a name for the folder and specify the folder attributes.




Figure 42 Create a New Folder page, General tab

3. In the **Compress** tab, indicate whether and how this folder and its contents are to be compressed, and if so, how.
4. After all information for the new folder is entered, click **OK**.

Deleting a folder

To delete a folder:

1. On the **Folders** page, navigate to the folder to delete. Select the folder, and then click **Delete**.

Summary information about the deletion is displayed.



NOTE:

View the summary information to confirm that this is the intended share.

2. Verify that the displayed folder is the folder to delete, and then click **OK**.

Modifying folder properties

To modify folder properties:

1. On the **Folders** page, navigate to the folder whose properties need to be edited, and then click **Properties**.




Figure 43 Folder Properties page, General tab

2. On the **General** tab, enter the new information for the folder.
3. On the **Compress** tab, indicate whether this folder and its contents are to be compressed, and if so, how.
4. After all changes have been completed, click **OK**.

Creating a new share for a volume or folder

Within the WebUI, there are two access points to the same screens used to create file shares:

- A share can be created for a folder while working with that folder in the **Folders** page.
- A share can be created and, if necessary, new folders can be created, while working with file shares on the **Shares** page.

This section discusses creating shares from the **Folders** page, and is an overview of the procedures. Complete details on the process of creating shares are included in the discussion that documents creating shares through the **Shares** tab. See the “[Managing shares](#)” section of this chapter for these details.



NOTE:

On select servers this function operates in a cluster but should only be used for non-cluster aware shares. Use Cluster Administrator to create shares for a cluster.

To create a new share for a specific volume or folder while on the **Folders** page:

1. Navigate to the desired volume or folder, and then click **Manage Shares**.
2. Click **New**.




Figure 44 Create New Share page, General tab

3. Enter the information for the share, including the name of the share, the allowed protocols, and corresponding permissions.



NOTE:

The Share path is the path of the previously selected volume or folder. This field is automatically completed by the system.

4. Select the appropriate tab to enter protocol specific information.

See the “[Managing shares](#)” section for detailed information about these entries.

5. After entering all share information, click **OK**.



NOTE:

The default permission settings for a new share are read-only.

Managing shares for a volume or folder

Within the WebUI, there are two access points to the same screens used to manage file shares:

- While working with a folder in the **Folders** pages, the administrator can create, delete, and modify shares for that folder.
- While working with file shares in the **Shares** pages, the administrator can create, delete, and modify shares (and if necessary, create new folders).



NOTE:

This section discusses managing shares from the Folders page, and is an overview of the procedures. Complete details on the process of managing shares are included in the discussion that documents creating shares through the Shares page. See the “[Managing shares](#)” section later in this chapter for these details.

To create, delete, and manage shares for a particular volume or folder while in the **Folders** page:

1. From the **Folders** directory, navigate to the target volume or folder, and then click **Manage Shares**.
All associated shares for that folder or volume are listed.
2. To create a new share, click **New**.
3. To delete a share, select the share to delete, and then click **Delete**.
4. To modify share properties, select the share to modify, and then click **Properties**.

Managing file level permissions

The WebUI of the storage server provides security at the share level, discussed later in this chapter. Security at the file level is managed using Windows Explorer available from the desktop of the storage server. To access the storage server desktop from the WebUI, go to the **Maintenance** tab, and then click **Remote Desktop**.

File level security includes settings for permissions, ownership, and auditing for individual files.

To enter file permissions:

1. Using Windows Explorer, navigate to the folder or file that needs to be changed and then right-click the folder.
2. Click **Properties**, and then click the **Security** tab.




Figure 45 Properties dialog box, Security tab

Several options are available on the **Security** tab:

- To add users and groups to the permissions list, click **Add**. Follow the dialog box instructions.
- To remove users and groups from the permissions list, highlight the desired user or group, and then click **Remove**.
- The center section of the **Security** tab provides a listing of permission levels. When new users or groups are added to the permissions list, select the appropriate boxes to configure the common file access levels.

3. To modify ownership of files or to modify individual file access level permissions, click **Advanced**.

Figure 46 illustrates the properties available on the **Advanced Security Settings** dialog box.




Figure 46 Advanced Security Settings dialog box, Permissions tab

Other functionality available in the **Advanced Security Settings** dialog box is illustrated in [Figure 46](#) and includes:

- **Add a new user or group**—Click **Add**, and then follow the dialog box instructions.
- **Remove a user or group**— Click **Remove**.
- **Replace permission entries on all child objects with entries shown here that apply to child objects**—This allows all child folders and files to inherit the current folder permissions by default.
- **Modify specific permissions assigned to a particular user or group**—Select the desired user or group, and then click **Edit**.

4. Enable or disable permissions by selecting the **Allow** box to enable permission or the **Deny** box to disable permission. If neither box is selected, permission is automatically disabled. [Figure 47](#) illustrates the **Edit** screen and some of the permissions.




Figure 47 User or Group Permission Entry dialog box

Another area of the **Advanced Security Settings** is the **Auditing** tab. Auditing allows you to set rules for the auditing of access, or attempted access, to files or folders. Users or groups can be added, deleted, viewed, or modified through the **Advanced Security Settings Auditing** tab.




Figure 48 Advanced Security Settings dialog box, Auditing tab

5. Click **Add** to display the Select User or Group dialog box.




Figure 49 Select User or Group dialog box



NOTE:

Click Advanced to search for users or groups.

6. Select the user or group.
7. Click **OK**.

The **Auditing Entry** dialog box is displayed.




Figure 50 Auditing Entry dialog box for folder name NTSF Test

8. Select the desired **Successful** and **Failed** audits for the user or group.
9. Click **OK**.



NOTE:

Auditing must be enabled to configure this information. Use the local Computer Policy Editor to configure the audit policy on the storage server.

The **Owner** tab allows taking ownership of files. Typically, administrators use this area to take ownership of files when the file ACL is incomplete or corrupt. By taking ownership, you gain access to the files, and then manually apply the appropriate security configurations.




Figure 51 Advanced Security Settings dialog box, Owner tab

The current owner of the file or folder is listed at the top of the screen. To take ownership:

1. Click the appropriate user or group in the **Change owner to** list.
2. If it is also necessary to take ownership of subfolders and files, enable the **Replace owner on subcontainers and objects** box.
3. Click **OK**.

Share management

There are several ways to set up and manage shares. The WebUI provides pages for setting up and managing shares. Additional methods include using a command line interface, Windows Explorer, or Management Console. This guide demonstrates using the WebUI to set up and manage shares.



NOTE:

Select servers can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses share setup for a non-clustered deployment. For information on managing file shares and printers in a cluster, see the “[Cluster Administration](#)” chapter.

As previously mentioned, the file sharing security model of the storage server is based on the NTFS file-level security model. Share security seamlessly integrates with file security. In addition to discussing share management, this section discusses share security. See “[Managing file level permissions](#)” earlier in this chapter for information on file security.

Share considerations

Planning the content, size, and distribution of shares on the storage server can improve performance, manageability, and ease of use.

The content of shares should be carefully chosen to avoid two common pitfalls: either having too many shares of a very specific nature or of having very few shares of a generic nature. For example, shares for general usage are easier to set up in the beginning, but can cause problems later. Frequently, a better approach is to create separate shares with a specific purpose or group of users in mind. However, creating too many shares also has its drawbacks. For example, if it is sufficient to create a single share for user home directories, create a "homes" share rather than creating separate shares for each user.

By keeping the number of shares and other resources low, the performance of the storage server is optimized. For example, instead of sharing out each individual user's home directory as its own share, share out the top level directory and let the users map personal drives to their own subdirectory.

Defining Access Control Lists

The Access Control List (ACL) contains the information that dictates which users and groups have access to a share, as well as the type of access that is permitted. Each share on an NTFS file system has one ACL with multiple associated user permissions. For example, an ACL can define that User1 has read and write access to a share, User2 has read only access, and User3 has no access to the share. The ACL also includes group access information that applies to every user in a configured group. ACLs are also referred to as permissions.

Integrating local file system security into Windows domain environments

ACLs include properties specific to users and groups from a particular workgroup server or domain environment. In a multidomain environment, user and group permissions from several domains can apply to files stored on the same device. Users and groups local to the storage server can be given access permissions to shares managed by the device. The domain name of the storage server supplies the context in which the user or group is understood. Permission configuration depends on the network and domain infrastructure where the server resides.

File-sharing protocols (except NFS) supply a user and group context for all connections over the network. (NFS supplies a machine-based context.) When new files are created by those users or machines, the appropriate ACLs are applied.

Configuration tools provide the ability to share permissions out to clients. These shared permissions are propagated into a file system ACL, and when new files are created over the network, the user creating the file becomes the file owner. In cases where a specific subdirectory of a share has different permissions from the share itself, the NTFS permissions on the subdirectory apply instead. This method results in a hierarchical security model where the network protocol permissions and the file permissions work together to provide appropriate security for shares on the device.



NOTE:

Share permissions and file level permissions are implemented separately. It is possible for files on a file system to have different permissions from those applied to a share. When this situation occurs, the file level permissions override the share permissions.

Comparing administrative (hidden) and standard shares

CIFS supports both administrative shares and standard shares.

- Administrative shares are shares with a last character of \$. Administrative shares are not included in the list of shares when a client browses for available shares on a CIFS server.
- Standard shares are shares that do not end in a \$ character. Standard shares are listed whenever a CIFS client browses for available shares on a CIFS server.

The storage server supports both administrative and standard CIFS shares. To create an administrative share, end the share name with the \$ character when setting up the share. Do not type a \$ character at the end of the share name when creating a standard share.

Planning for compatibility between file sharing protocols

When planning for cross-platform share management on the storage server, it is important to understand the different protocols and their associated constraints. Each additional protocol that is supported adds another level of constraints and complexity.

NFS compatibility issues

When planning to manage CIFS and NFS shares, consider two specific requirements.



NOTE:

Further information, including details about the NFS Service and the User Mapping service, is available in the “[Services for NFS/UNIX](#)” chapter.

- **NFS service does not support spaces in the names for NFS file shares.**
NFS translates any spaces in an export into an underscore character. Additional translations can be set up for files. (See the “OEM Supplemental Help” chapter of the SFU help, found on the storage server). This feature ensures the greatest level of compatibility with NFS clients, because some do not work with NFS exports that contain a space in the export name.

If you plan to use the same name when sharing a folder through CIFS, and then exporting it through NFS, do not put spaces in the CIFS share name.
- **NFS service does not support exporting a child folder when its parent folder has already been exported.**
An NFS client can access a child folder by selecting the parent folder and then navigating to the child folder. If strict cross-platform compatibility is an administration goal, CIFS must be managed in the same way. Do not share a folder through CIFS if the parent folder is already shared.

Managing shares

Shares can be managed through the **Shares** tab of the WebUI. Tasks include:

- Creating a new share
- Deleting a share
- Modifying share properties
- Publishing in DFS (see “[Publishing a new share in DFS](#)”)



NOTE:

These functions can operate in a cluster on select servers but should only be used for non-cluster aware shares. Use Cluster Administrator to manage shares for a cluster. The page will display cluster share resources.

Creating a new share

To create a new share:

1. From the WebUI, click the **Shares** tab, and then click **Shares**.
2. Click **New**.

The screenshot shows the 'Shares' tab selected in the top navigation bar. Below it, the 'New Share' dialog is open. On the left, there's a sidebar with tabs for 'General', 'Windows Sharing', 'UNIX Sharing', and 'Web Sharing'. The 'General' tab is active. In the main area, there are fields for 'Share name' and 'Share path', both currently empty. A checkbox labeled 'Create folder' is unchecked. Below these, there's a section titled 'Select the clients for which you want to allow access to the share.' with three checkboxes: 'Windows (Microsoft SMB)' (checked), 'UNIX (NFS)' (unchecked), and 'Web (HTTP)' (unchecked, with a note '- Web sharing is currently stopped.'). Further down, there's a note about publishing to a DFS root and a checkbox 'Publish to DFS root' which is unchecked. At the bottom right of the dialog are 'OK' and 'Cancel' buttons.

Figure 52 Create a New Share page, General tab

3. On the **General** tab, enter the following information:
 - Share name
 - Share path
 - Client protocol types

To create a folder for the new share, select the indicated box. The system creates the folder at the same time it creates the share.

Protocol specific tabs are available to enter sharing and permissions information for each sharing type. See “[Modifying share properties](#)” for detailed information on these tabs.

4. After entering all share information, click **OK**.

Deleting a share



CAUTION:

Before deleting a share, warn all users to exit that share and confirm that no one is using the share.

To delete a share:

1. From the **Shares** menu, click **Delete**.
2. Verify that this is the correct share, and then click **OK**.

Modifying share properties

To change share settings:

1. From the **Shares** menu, select the share to modify, and then click **Properties**.




Figure 53 Share Properties page, General tab

2. Select the appropriate box to select the protocol to use. To change client protocol information, click the corresponding tabs.
 - Windows Sharing
 - UNIX Sharing
 - Web Sharing (HTTP)
3. After all share information has been entered, click **OK**.

Windows sharing

From the **Windows Sharing** tab of the **Share Properties** page:

1. Enter a descriptive **Comment**, and the **User limit** (both optional).




Figure 54 Share Properties page, Windows Sharing tab

2. Select Offline settings.
3. Set the permissions.

The **Permissions** box lists the currently approved users for this share.

- To add a new user or group, either select a user or group from the box at the bottom right of the screen or manually enter the user or group name in the **Add a user or group** box, and then click **Add**. That user or group is added to the Permissions box.
- To remove access to a currently approved user or group, select the user or group from the Permissions box, and then click **Remove**.
- To indicate the type of access allowed for each user, select the user, and then expand the Allow and Deny drop down boxes. Select the appropriate option.

4. After all Windows Sharing information is entered, either click the next **Sharing** tab or click **OK**.

UNIX sharing

From the **UNIX Sharing** tab of the **Share Properties** page:

1. Indicate the machines that will have access to this share.

Select the machine to include in the **Select a group** box or manually enter the NFS client computer name or IP address, and then click **Add**.




Figure 55 Share Properties page, UNIX Sharing tab

2. Indicate the access permissions.

Select the machine from the Permissions list, and then select the appropriate access method from the **Access permissions** drop down box.

The types of access are:

- **Read-only**— Use this permission to restrict write access to the share.
- **Read-write**— Use this permission to allow clients to read or write to the share.
- **No access**— Use this permission to restrict all access to the share.

3. Select whether or not to allow root access.

- **Read-only + Root**— Use this permission to restrict write access to the share. Use this permission to assign administrative access to the share. This will allow the client computer to have root access to the NFS share. Map the UNIX root user to the Windows user Administrator. Also, map the group that this UNIX root belongs to, to the Windows group Administrator.
- **Read-write + Root**— Use this permission to allow clients to read or write to the share. Use this permission to assign administrative access to the share. This will allow the client computer to have root access to the NFS share. Map the UNIX root user to the Windows user Administrator. Also, map the group that this UNIX root belongs to, to the Windows group Administrator.

4. After all UNIX sharing information is entered, click **OK**.

Web sharing (HTTP)

From the **Web Sharing** tab of the **Share Properties** page:

1. Select the read and write access permissions that are allowed.
2. Click **OK**.

AFP (Appletalk) sharing

AppleTalk shares can be set up only after AppleTalk Protocol and File Services for Macintosh have been installed on the storage server.



NOTE:

AppleTalk shares should not be created on clustered resources as data loss can occur due to local memory use.

Installing the AppleTalk Protocol

To install the AppleTalk Protocol:

1. From the desktop of the storage server, select **Start > Settings > Network Connections**. Right-click **Local Area Connection**, and then click **Properties**.
2. Click **Install**.




Figure 56 Local Area Connection Properties page, Install option

3. Select **Protocol**, and then click **Add**.
4. Select **AppleTalk Protocol**, and then click **OK**.

Installing File Services for Macintosh

To install File Services for Macintosh:

1. From the WebUI, click the **Maintenance** tab.
2. Click **Remote Desktop**.
3. Open **Add/Remove Programs** from the Control Panel.
4. Click **Add/Remove Windows Components**.
5. Double-click **Other Network File and Print Services**.
6. Select **File Services for Macintosh**, and then click **OK**.
7. Click **Next**.

8. Click **Finish**.

Setting AppleTalk Protocol Properties

To set AppleTalk Protocol properties:

1. From the WebUI, click the **Shares** tab.
2. Click **Sharing Protocols**.
3. Click the AppleTalk radio button, and then choose **Properties**.
4. Insert login message, if desired.
5. Under Security, “Enable client authentication with,” select Apple Clear Text or Microsoft.

To set up AppleTalk shares, from the WebUI:

1. Click **Shares**.
2. Click **Shares** again.
3. Click **New**.
4. Enter the share name and share path.
5. Select Apple Macintosh. Clear other file types if necessary.
6. Click **AppleTalk Sharing**.
7. Enter a user limit.
8. Enter password information.
9. Indicate whether the share has read only permission or read write permission.
10. After all AppleTalk Sharing information is entered, click **OK**.

Protocol parameter settings

As previously mentioned, the storage server supports the following protocols:

- DFS
- NFS
- FTP
- HTTP
- Microsoft SMB

This section discusses the parameter settings for each protocol type.



NOTE:

See the Protocol section of the Cluster Administration chapter for information about protocol selection and management in a cluster.

To access and enter protocol parameter settings:

1. From the WebUI, click the **Shares** tab.
2. Click **Sharing Protocols**.

Name	Status	Startup Type	Description	Tasks
DFS	Running	Automatic	Allows organization of distributed data shares into a logical namespace for simplified data access	<input type="button" value="Enable"/> <input type="button" value="Disable"/> <input type="button" value="Properties..."/>
FTP	Stopped	Disabled	Allows access to data shares from FTP clients	
HTTP	Running	Automatic	Allows access to data shares from Web browsers	
Microsoft SMB	Running	Automatic	Allows access to data shares from Microsoft Windows clients	
NFS	Running	Automatic	Allows access to data shares from UNIX network file system clients	

Figure 57 File Sharing Protocols page

3. Protocols and their statuses are listed. The following options are available:
 - Enabling a protocol
 - Disabling a protocol
 - Modifying Protocol Settings

Because enabling and disabling a protocol are self explanatory, only modifying protocol specific settings is described in this section.

DFS protocol settings

Using Distributed File System (DFS), system administrators can make it easy for users to access and manage files that are physically distributed across a network. Users do not need to know and specify the actual physical location of files in order to access them.

For example, if documents are scattered across multiple servers in a domain, DFS can make it appear as though the documents all resides on a single server. This eliminates the need for users to go to multiple locations on the network to find the information.

Each DFS namespace requires a root. A DFS root is a starting point of the DFS namespace. The root is often used to refer to the namespace as a whole. A root maps to one or more root targets, each of which corresponds to a shared folder on a server. A root is implemented as a shared folder on the DFS server.

Deploying DFS

A distributed file system can be implemented as a stand-alone root distributed file system or as a domain root distributed file system. The type of a distributed file system determines which client computers can access the distributed file system.

A stand-alone DFS root:

- Does not use Active Directory to manage DFS.
- Cannot have more than one root on a server.
- Does not support automatic file replication using the File Replication service (FRS).
- Is not fault tolerant and if the root fails the entire namespace will collapse.


A domain DFS root:

- Must be hosted on a domain member server.
- Has its DFS namespace automatically published to Active Directory.
- Can have more than one root on a server.
- Supports automatic file replication through FRS.
- Supports fault tolerance through FRS.

Two points of management of the DFS namespace are provided with the storage server. These points of management are the WebUI and the Distributed File System Administration Tool located on the local console of the storage server under **Start > Programs > Administrative Tool**. (See [Figure 58](#)). The WebUI is designed to provide the following functions:

- Stand alone root management (Add, Delete)
- Share publishing to stand alone or domain DFS
- Default behavior for DFS share publishing

All other functions must be performed via the DFS Administration Tool.



[Figure 58 DFS Win32 GUI](#)

DFS Administration Tool

The DFS Administration Tool provides extended functionality not found in the WebUI. These functions include:

- Management of multiple DFS Roots on multiple machines from a single interface
- Domain based DFS management
- Target and Link management
- Status Checks of a DFS managed share link
- Exporting of the DFS names space to a text file

The storage server administration guide only provides instructions on the WebUI portion of the product. The DFS Administration Tool is complete with online help. In addition, general information on DFS may be found at:

<http://www.microsoft.com/windowsserver2003/techinfo/overview/dfs.mspx>

Accessing the DFS namespace from other computers

In addition to the server-based DFS component of the Windows Storage Server 2003 family, there is a client-based DFS component. The DFS client caches a referral to a DFS root or a DFS link for a specific length of time, defined by the administrator.

The DFS client component runs on a number of different Windows platforms. In the case of older versions of Windows, the client software must be downloaded to run on that version of Windows. Newer versions of Windows have client software built-in.

Non-Windows (such as Linux/UNIX) based clients can not access the DFS namespace as DFS is dependent on a Windows component to function.

Setting DFS sharing defaults

The WebUI can be used to set the default DFS settings provided when creating a shared folder. When a new shared folder is created, the DFS defaults may be overridden.

To set DFS sharing defaults:

1. From the WebUI, click the **Shares** tab.
2. Click **Sharing Protocols**.
3. Click **DFS**, and then click **Properties**.




Figure 59 DFS Properties page, General tab

4. On the **General** tab, select the default settings that are desired when creating a shared directory.
 - To set the default to publish the share to the local DFS root, select **Publish new shares to the local DFS root**.
 - To set the default to publish the share to another DFS root, select **Publish new shares to a domain DFS root or a stand-alone DFS root on a different server**. In the DFS root box, type the path of the default DFS root.
 - To not publish the share to a DFS root, select **Do not publish new shares to a DFS root**.
5. Click **OK**.

Creating a local DFS root

The WebUI can be only be used to create a single, local stand-alone DFS root on the server as mentioned previous. To create a domain DFS root use the DFS administrative tool. For more information about DFS root types refer to the section above entitled “[Deploying DFS](#).”

To create a local stand-alone DFS root:

1. From the WebUI, click **Shares**.
2. Click **Sharing Protocols**.
3. Click **DFS**, and then click **Properties**.




Figure 60 DFS Properties page, Local DFS Root tab

4. On the **Local DFS Root** tab, enter the name of the DFS root in the **Root name** box.
5. In the **Folder path** box, enter the path of the folder that corresponds to the root. Click **Create folder** if the folder does not exist.
6. Click **Create DFS Root**, and then click **OK**.

Deleting a local DFS root

The WebUI enables the deletion of a local stand-alone DFS root on the server only. The DFS Administrative Tool must be used to manage Domain DFS Roots. Hence, if there is more than one root on the server, the first root (in alphabetical order, with local stand-alone roots grouped ahead of domain roots) is available to be deleted. If only domain roots exist on the server, the first domain root is listed, but it cannot be deleted using the WebUI. The WebUI can only be used to manage local stand-alone DFS roots.

To delete a local DFS root:

1. From the WebUI, click the **Shares** tab.
2. Click **Sharing Protocols**.
3. Click **DFS**, and then click **Properties**. On the **Local DFS Root** tab, click **Delete Root**.
4. Click **OK**.

Publishing a new share in DFS

After a root has been established either on the local machine or one in the network, shares can be published to extend the virtual name space. For example, several shares can be created for a DFS root labeled "Engineering." The shares might be titled "Documentation," "Technical Specs," and "Project Info." When mapping to `\computername\engineering`, all three of these shares would be found under the

mapped drive even though they exist on different storage servers, drives, or shares points. To publish a share in a DFS root:




Figure 61 DFS share example

1. From the WebUI, click the **Shares** tab.
2. Enter a new share name.
3. Enter a folder name (select the checkbox **Create folder** if appropriate).
4. Verify that the Windows checkbox is selected. (DFS is dependent on the SMB protocol).
5. Under DFS, select the box if unchecked.



NOTE:

The default behavior can be set to publish all shares to DFS. In this case the box will be checked. See the section above Setting DFS Sharing Defaults.

6. Enter the name of the DFS root to publish the share ("Engineering" in this example). The network name is displayed below the entry.
7. Click **OK**.

A share name is published in the namespace.

To view the namespace, map a drive to the DFS root. All published shares are seen in the namespace.




Figure 62 DFS share example, mapped drive

In this case, *Documentation* exists on *G:\documentation*, *Technical Specs* exists on *G:\technical specs* and *Project Info* exists on *C:\project info* on the local machine but they are all accessible via *\\\DOCBOX\engineering*.

Publishing an existing share in DFS

To enable an existing shares for DFS:

1. From the WebUI, click the **Shares** tab.
2. Select the target share from the table, and then click **Publish in DFS**.
3. Enter the name of the DFS root to publish the share to.
4. Click **OK**.

The share appears in the DFS underneath the DFS root.

Removing a published share from DFS

After a share is published in DFS, it can be removed from the virtual namespace via the **Shares Property** page. To remove a share from DFS perform the following steps:

1. From the WebUI, click the **Shares** tab.
2. Select the target share from the table and select properties.

3. Clear the **Publish to DFS root** box.

4. Click **OK**.

The share no longer appears in the DFS.

Storage management



NOTE:

The storage management features are not offered with all ProLiant Storage Server models. See the user guide for a listing of supported models.

The storage management features built into the storage server are composed of three main features, and are applicable at the directory level of a share. These features include:

- Directory Quotas
- File Screening
- Storage Reports

For procedures and methods, refer to the online help available within the WebUI via the ? in the right hand corner of each accompanying feature management page.



NOTE:

The storage management features are not supported in a clustered environment. In a clustered environment these features should be uninstalled as instructed in the Cluster Installation Guide on select servers. (See [Figure 63](#)).




Figure 63 Uninstall storage manager

Directory quotas

Directory quotas provide a way to limit and monitor space consumed by all files in a folder. For information on setting quotas on volumes, see Chapter 3.

Directory quotas limit the size of the managed object regardless of who writes to, or who owns files in the managed object. For example, if a 50 MB directory quota is set on the managed object `c:\users\JDoe`, that directory and all its contents is limited to 50 MB regardless of who owns the files in that directory or who writes files to that directory.

Directory quotas allow for the addition, deletion, monitoring, and changing of space limits for selected directories on the storage server. Directory quotas provide disk space monitoring and control in real time, and support active and passive limits with two real-time space alarms.

The Directory Quota feature includes the following components:

- Active and passive space limits on directories
- Best practice storage resource management policies
- Severe alarm threshold
- Warning alarm threshold
- Auto discovery of drives
- Customized messages
- Alarms sent to the event log
- Alarms sent to the user
- Storage reports that can be sent to an intranet web site
- Custom script support

The directory quota set on the system partition always has a passive limit and uses device size (capacity). If the system does not have sufficient quota to write files, it may fail. Also, if the system partition does not have enough space to write temporary files during boot, the system may not restart. Avoid this by using caution when placing quotas on the system directories.

Directory quotas use each file's allocation size to determine how much space is used. The allocation size is slightly larger than the actual space used as displayed by Windows Explorer and other Windows programs for the data in a file. This discrepancy may cause some confusion, but the Directory Quota feature is correctly charging the user for the amount of disk space actually consumed to store a file. Large cluster sizes on file allocation table (FAT) file systems may add to the confusion because the entire cluster is always allocated, regardless of the file size. NTFS file systems store very small files in the index file and typically have more reasonable cluster sizes.

Because of the differences in the amount of storage requested for a file extension operation and the amount actually allocated by Windows Storage Server 2003 for that extension, the user may be allowed to exceed his quota by as much as one cluster. For example, assume the user has a quota of 100 KB and has used 96 KB on a file system with a cluster size of 8 KB. The user creates a file of 1 KB. Windows Storage Server 2003 requests 1024 bytes be allocated for the file. Since this is less than the remaining quota for the user, the operation is allowed to continue. However, if the cluster size is 8 KB, Windows Storage Server 2003 will actually allocate 8 KB for the file. The user has now used 104 KB, and while this is allowed, future attempts to create or extend files will fail.

Establishing directory quotas

Directory quotas are established in a two-part fashion. First, a policy is defined using the policies selection from the Directories Policy Page. After a policy is established it can be assigned to a particular directory via the WebUI "New Directory Quota Wizard." By default there are a number of predefined policies, which include:

- 100 MB limit
- 500 MB limit
- Best Practices report
- Default
- Monitor directory
- Partition alert

Each of these policies provides an example of a particular policy type. Custom policies should be created to meet the needs of the environment.

The screenshot shows the 'Shares' section of the administration interface. The 'Shares' tab is selected at the top. Below it, the 'Directory Quota' tab is active. A sub-section titled 'Directory Quota Policies' is displayed. The main area lists several policies with checkboxes next to their names: '100 MB Limit', '500 MB Limit', 'Best Practice Reports', 'Default', 'Monitor Directory', and 'Partition Alert'. To the right of this list is a vertical 'Tasks' column containing three buttons: 'New...', 'Delete', and 'Properties...'. At the bottom left of the main area is a 'Back' button.

Figure 64 Directory Quota Policies page

Within each policy, there are a number of configuration screens that are presented in the form of a wizard. The wizard collects the following information to create a policy:

- Policy name
- Disk space limit and unit of measurement
- Passive limit (If selected the limit issues warnings but does not prevent access.)
- Alarm threshold for severe and warning messages
- Notification for severe and warning messages

The notification field allows the creation of a message to be sent to the eventlog of the server or via Netbios as a popup on the client machine. Netbios is not supported in all customer environments and the popup function may not be supported. The notification includes macro functionality and variable definitions for user custom messages. The help function (?) in the right hand corner of the UI provides an online guide to building these macro function messages under the topic "Directory Quota Alarm Notification."

To modify any of these settings at a later time, click **Properties** after selecting a particular policy or quota. In addition to policy settings for existing shares, default policies can be set in advance for new devices added to the system by clicking **Preferences** on the **Directory Quota** page.

File screening

File screening allows the administrator to limit or monitor files based on extension, for example disallow all .pst and .mp3 files. The filter is merely based on extensions and not the content of the files. Hence, if a file extension is renamed away from .mp3 for example to .mpp, the filter software allows the file to be stored. A complete online help guide in the WebUI is provided for file screening via the ? in the right hand corner of the UI.

File screening is established in the policy settings. Screening groups contain a collection of authorized and un-authorized file extensions. Filters determine which folders to exclude. Alarms, similar to the actions when a quota threshold is exceeded, can be set up when an unauthorized file type is set up.

File screening includes the following features:

- Active and passive file screening on directories
- Best practice file screening policies
- Notification alarm when file screening policy is violated
- Audit database containing screened files
- Customized alarm messages
- Alarm messages to the event log
- Alarm messages to a user
- Storage reports when alarm is activated and sent to intranet web site
- Custom script when alarm is activated
- Real time monitoring of file screening activity

Use caution when placing screening parameters on the system partition. If certain classes of files are screened from the system partition, the operating system may not have the access to save temporary working files. It is a good idea to exclude systems directories from screening. Another option is to create a passive screening policy that allows files to be saved, but the file activity to be logged.

File Screening essentially has the same feature sets as directory quotas with one exception: groupings of file types, such as multimedia files, graphics, and so on, are created first. These groups are then placed in a particular policy. A file screen is then enabled on a directory and the various policies are applied for a particular directory. Lastly, the same types of alert notification are allowed as in the case of the directory quotas. See the online help for additional information.

Storage reports

Storage reports allow the administrator to analyze the contents of the storage server via standard reports for common tasks. The reports can be displayed using text, simple HTML tables, or Active HTML. When using Active HTML, the ActiveX control provides graphs. A complete online help guide in the WebUI is provided for reporting via the ? in the right hand corner of the UI.

Reports can be scheduled, or produced on demand.

Storage reports address disk usage, wasted space, file ownership, security, and administration. Reports can run interactively, be scheduled on a regular basis, or run as part of a storage resource management policy when disk space utilization reaches a critical level.

Storage reports may be presented in HTML and text (.txt) formats. The output formats can be e-mailed to a list of users.

The following features are included with storage reports:

- Best practice storage resource management reports
- Integration with best practice storage resource management policies
- Scheduled storage reports
- Storage reports sent to an intranet web site
- Storage reports sent to users through e-mail



NOTE:

Large storage reports should be scheduled for off-hours.

Print services (where licensed)



NOTE:

Print services are not available on all models. Refer to the installation guide to determine availability of this feature.

Printer services support network printers only and are not intended for use with locally attached printers (USB or Parallel port connected).



NOTE:

See the [Cluster Administration](#) chapter for information on clustering a print spooler.

If the storage server is a part of an Active Directory domain rather than Workgroup, the storage server enables the following management features:

- Restrict access to a printer based domain user accounts
- Publish shared printers to Active Directory to aid in search for the resource

Before adding a print server role, the following checklist of items should be followed:

1. **Determine the operating system version of the clients that will send jobs to this printer.** This information is used to select the correct client printer drivers for the client and server computers utilizing the printer. Enabling this role on the print server allows the automatic distribution of these drivers to the clients. Additionally, the set of client operating systems determines which of these drivers need to be installed on the server during the print server role installation.
2. **At the printer, print a configuration or test page that includes manufacturer, model, language, and installed options.** This information is needed to choose the correct printer driver. The manufacturer and model are usually enough to uniquely identify the printer and its language. However, some printers support multiple languages, and the configuration printout usually lists them. Also, the configuration printout often lists installed options, such as extra memory, paper trays, envelope feeders, and duplex units.
3. **Choose a printer name.** Users running Windows-based client computers choose a printer by using the printer name. The wizard that you will use to configure your print server provides a default name, consisting of the printer manufacturer and model. The printer name is usually less than 31 characters in length.
4. **Choose a share name.** A user can connect to a shared printer by entering this name, or by selecting it from a list of share names. The share name is usually less than 8 characters for compatibility with MS-DOS and Windows 3.x clients.
5. (Optional) **Choose a location description and a comment.** These can help identify the location of the printer and provide additional information. For example, the location could be "Second floor, copy room" and the comment could be "Additional toner cartridges are available in the supply room on floor 1."

Configuring the print server

To set up a print server:

1. Select **Start > Settings > Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.
2. Click **Add or Remove a Roll**.

A wizard starts.

3. Click **Next**.
4. Select Printer Server in the list of Server Roles, and then click **Next**.
5. Select Windows 2000 and Windows XP clients only, and then click **Next**.



NOTE:

While the "All Windows" support may be selected at this step, it is more efficient to add the alternative operating systems on each printer after the wizards are complete. See "[Adding additional operating system support](#)."

6. Click **Next** on the **Summary** page to start the Add Printer Wizard.
7. Select Local Printer, clear "automatically detect install my plug and play printers," and then click **Next**.



NOTE:

Local Printer is used to create a TCP/IP port connection to a network enabled printer over the network. The storage server only supports network attached printers and does not support directly connected printers via USB or Parallel Port.

8. Select **Create a new port**, and then select **Standard TCP/IP Port** (recommended).

The Add Standard TCP/IP Printer Port Wizard starts.

9. Click **Next**.
10. Enter the name or IP address of the printer. The IP address is usually listed on the printer configuration page. The wizard completes the Port Name field. Click **Next**.
11. The wizard attempts to connect to the printer. If the wizard is able to connect, the **Completing the Add Standard TCP/IP Printer Port** Wizard page opens; click **Finish**. If the wizard is not able to connect, the **Additional Port Information Required** page opens.
 - a. Verify that the IP address or name is correct.
 - b. Select **Standard** to identify the printer network adapter. A list of manufacturers and models of the network adapters is displayed. Select the appropriate printer in the Standard list.
 - c. If the printer network adapter uses nonstandard settings, click **Custom**, and then click **Settings**. The **Configure Standard TCP/IP Port Monitor** page opens. Specify the settings that are recommended by the manufacturer of the printer network adapter, and then click **OK**.
 - d. Click **Next**.

12. Select the manufacturer and the type of printer in the presented list, and then click **Next**. If the printer does not exist in the list, click **Have disk** and load the drivers, or select a compatible driver.



NOTE:

On models having print services available, additional print drivers are located on the storage server in the *c:\hpnas\components\HP Storage Server Printer Drivers* folder.

13. Enter the name of the desired printer to be presented on the storage server, and then click **Next**.
14. Enter a Share Name for the printer to be used on the network, and then click **Next**.
15. Enter a location description and a comment, and then click **Next**.
16. Select **Print a test page**, and then click **Next**.
17. Clear the **Restart the add printer wizard** if adding only one printer, and then click **Finish**.
A test page prints.
18. Click **OK** if the page printed, otherwise click **Troubleshoot**.
If **Restart the add printer wizard** was selected, the wizard restarts to add an additional printer.
19. Repeat the steps above for adding an additional printer.

Removing the print server role

To remove the print server role:

1. Select **Start > Settings > Control Panel**, double-click **Administrative Tools**, and then double-click **Manage Your Server**.
2. Click **Add or Remove a Roll**.
A wizard starts.
3. Click **Next**.
4. Select **Printer Server** in the list of Server Roles, and then click **Next**.
5. Select the checkbox **Remove the printer role**, and then click **Next**.
The printer role is removed.
6. Click **Finish**.

Adding an additional printer

To add additional printers to the storage server:

1. Select **Start > Settings > Printers and Faxes > Add Printer**.
The add printer wizard starts.
2. Click **Next**.
3. Select **Local Printer** and clear **Automatically detect install my plug and play printers**.

4. Click **Next**.



NOTE:

Local Printer is used to create a TCP/IP port connection to a network enabled printer over the network. The storage server only supports network attached printers and does not support directly connected printers via USB or Parallel Port.

5. Click **Create a new port**, and then select **Standard TCP/IP Port** (recommended).

The **Add Standard TCP/IP Printer Port** wizard starts.

6. Click **Next**.

7. Enter the name or IP address of the printer. The IP address is usually listed on the printer configuration page. The wizard completes the Port Name field. Click **Next**.

8. The wizard attempts to connect to the printer.

- If the wizard is able to connect, the **Completing the Add Standard TCP/IP Printer Port** wizard page opens. Click **Finish**.
- If the wizard is not able to connect, the **Additional Port Information Required** page opens:
 - Verify that the IP address or name that was entered is correct.
 - Select **Standard** to identify the printer network adapter. A list of manufacturers and models of the network adapters is displayed. Select the appropriate printer from the Standard list.
 - If the printer network adapter uses nonstandard settings, click **Custom**, and then click **Settings**. The **Configure Standard TCP/IP Port Monitor** page appears. Specify the settings that are recommended by the manufacturer of the printer network adapter, and then click **OK**.
 - Click **Next**.

9. Select the manufacturer and the type of printer in the presented list, and then click **Next**. If the printer does not exist in the list, click **have disk** and load the drivers or select a compatible driver.



NOTE:

On models having print services available, additional print drivers are located on the storage server in the *c:\hpnas\components\HP Storage Server Print Drivers* folder.

10. Enter the name of the desired printer to be presented on the storage server, and then click **Next**.

11. Enter a Share Name for the printer that will be used on the network, and then click **Next**.

12. Enter a location description and a comment, and then click **Next**.

13. Select **Print a test page**, and then click **Next**.

14. Click **Finish**.

A test page prints.

15. Click **OK** if the page printed, otherwise click **Troubleshoot**.

Adding additional operating system support

By default, support is added for Windows 2000 and Windows XP. If the client base is composed of other Windows operating systems, additional printer drivers must be loaded. To load an additional driver for client download:

1. Select **Start > Settings > Printers and Faxes**, and then right-click the printer to manage.
2. Click **Properties**.
3. Click the **Sharing** tab.
4. Click **Additional Drivers**.
5. Click the desired operating systems, and then click **OK**.
6. A dialog box opens to add the additional drivers from disk.

Installing print services for UNIX

1. Log on as administrator or as a member of the Administrators group.
2. Select **Start > Control Panel**, and then click **Add or Remove Programs**.
3. Click **Add/Remove Windows Components**.
4. In the **Components** list, click **Other Network File and Print Services** (but do not select or clear the check box), and then click **Details**.
5. In the Subcomponents of **Other Network File and Print Services** list, select **Print Services for UNIX**, if appropriate to the print services that you want to install:

Print Services for UNIX: This option permits UNIX clients to print to any printer that is available to the print server.



NOTE:

When installing Print Services for UNIX, this automatically installs the LPR port and the TCP/IP Print Server service.

6. Click **OK**, and then click **Next**.
7. Click **Finish**.

HP Web Jetadmin

HP Web Jetadmin is a simple peripheral management software for remotely installing, configuring, and managing a wide variety of HP and non-HP network peripherals using only a standard web browser. The following URL provides additional feature information, plus a link to download the software:

http://h10010.www1.hp.com/wwpc-JAVA/offweb/vac/us/en/en/network_software/wja_overview.html

7 Services for NFS/UNIX

Microsoft Services for NFS and Windows Services for UNIX are comprehensive software packages designed to provide complete UNIX environment integration into a Windows NT, Windows 2000, Windows Storage Server 2003, or Active Directory domain file server. Services for NFS manages tasks on both Windows and UNIX platforms. Tasks include creating NFS exports from Windows and administering user name mappings.



NOTE:

To determine which version is on your storage server, select **Start > Programs**. The program list displays either Microsoft Services for NFS or Windows Services for UNIX.

The following Services for NFS components are included in the storage server:

- Server for NFS
- User Name Mapping
- NFS Authentication
- SFU 3.5



NOTE:

Services for NFS/UNIX can be implemented in both clustered and non-clustered environments using select storage servers. This chapter discusses Services for NFS/UNIX in a non-clustered deployment. For additional information that is specific to a cluster, see the Cluster Administration chapter.

Server for NFS

Services for NFS/UNIX enables UNIX clients to access a file share on the storage server. The Services for NFS server supports NFS Version 2 and Version 3, and supports them both on the TCP and UDP network protocols.

Services for NFS/UNIX is more fully integrated into the operating system than other third-party NFS server packages. The administrative interface for NFS exports is similar to the Server Message Block (SMB) sharing interface used by Windows platforms. With Server for NFS properly configured, the administrator can create shares that are simultaneously accessible by multiple client types. For example, some of the options for shares include configurations for CIFS/SMB sharing only, simultaneous NFS/CIFS/SMB sharing, simultaneous NFS/CIFS/SMB/HTTP sharing, or simply NFS only sharing.

Authenticating user access

NFS export access is granted or denied to clients based on client name or IP address. The server determines whether a specific client machine has access to an NFS export. No user logon to the NFS server takes place when a file system is exported by the NFS server. Permission to read or write to the

export is granted to specific client machines. For example, if client machine M1 is granted access to an export but client M2 is not, user jdoe can access the export from M1 but not from M2.

Permissions are granted on a per-export basis; each export has its own permissions, independent of other exports on the system. For example, file system a can be exported to allow only the Accounting department access, and file system m can be exported allowing only the Management department access. If a user in Management needs access to the Accounting information, the A export permissions can be modified to let that one user's client machine have access. This modification does not affect other client access to the same export, nor does it allow the Management user or client access to other exports.

After the client machine has permission to the export, the user logon affects file access. The client machine presents the UNIX user ID (UID) and group ID (GID) to the server. When the computer accesses a file, the UID and GID of the client are transferred to a Windows user ID and group ID by the mapping server. The ACLs of the file or directory object being requested are then compared against the mapped Windows login or group ID to determine whether the access attempt should be granted.



NOTE:

User credentials are not questioned or verified by the NFS server. The server accepts the presented credentials as valid and correct.

If the NFS server does not have a corresponding UID or GID, or if the administrator has set other conditions to filter out the user, a process called squashing takes effect. Squashing is the conversion of an unknown or filtered user to an anonymous user. This anonymous user has very restricted permissions on the system. Squashing helps administrators manage access to their exports by allowing them to restrict access to certain individuals or groups, and to squash all others down to restricted (or no) access. Squashing enables the administrator to allow permissions instead of denying access to all the individuals who are not supposed to have access. See "NFS User and Group Mappings" later in this chapter for specific information about creating and maintaining mappings.

S4U2 functionality

Windows Server 2003 Active Directory now has support for the S4U2Proxy extension to the Kerberos protocol. This extension allows services in the domain to act on behalf of a user. Therefore, you do not have to install the Server for NFS Authentication dll on domain controllers on a total Windows Server 2003 domain for Server for NFS to authenticate domain users. For more information on the S4U2Proxy, consult the S4U2Self topic in the following URL:

<http://msdn.microsoft.com/msdnmag/issues/03/04/SecurityBriefs/default.aspx>



NOTE:

The S4U2 functionality does not work until the domain functional level is elevated to Windows Server 2003.

To elevate the functional level to Windows Server 2003:

1. On the Windows 2003 domain controller, open Active Directory Domains and Trusts.
2. In the console tree, right-click the domain for which you want to raise functionality, and then click Raise Domain Functional Level.
3. In Select an available domain functional level, click **Windows Server 2003**.

4. Click **Raise**.

NFS Authentication is still the primary user name mapping authentication method used for domain mappings. If NFS Authentication fails it will try to use S4U2. Thus, the NFS Authentication dll is still the primary method with S4U2 being the backup method.

Indicating the computer to use for the NFS user mapping server

During the processes of starting and installing the storage server, the name localhost is assigned by default to the computer. It is assumed that the storage server is the computer that will be used for user name mapping.

If there are other mapping servers and a machine other than the localhost that will store user name mappings, the name of that computer must be indicated, as detailed below:

1. Use **Remote Desktop** to access the **Management Console**, click **File Sharing, Microsoft Services for Network File System**. Click **Settings**. [Figure 65](#) is an example of the Server for NFS user interface.
2. In the **Computer** name box of the user-mapping screen, type the name of the computer designated for user mapping and authentication.
3. Localhost is the computer name assigned by default on the storage server. To control user mapping from a different computer, enter the name of that computer.



NOTE:

If a machine other than the localhost is to be used, make sure that the user name mapping service is installed and running on that machine.



NOTE:

If the authentication software is not installed on all domain controllers that have user name mappings, including Primary Domain Controllers, Backup Domain Controllers, and Active Directory Domains, then domain user name mappings will not work correctly.




Figure 65 Microsoft Services for NFS screen, Settings tab

Logging events

Various levels of auditing are available. Auditing sends Services for NFS events to a file for later review and establishes log-setting behavior. Some behavior examples include events logged and log file size. See the online Services for NFS help for more information.

1. Use Remote Desktop to access the Management Console, click **File Sharing, Microsoft Services for Network File System, Server for NFS**. Click the **Logging** tab.
2. To log events to the event viewer application log, click the check box for **Log events to event log**.
3. To log selected event types, select the check box for **Log events in this file** on the screen.
4. Enter a filename, or use the default filename provided, and log file size (7 MB default). The default log file is created when the changes are applied.




Figure 66 Server for NFS screen, Logging tab

Server for NFS server settings

The storage server has new features for Services for NFS included in the Services for NFS administration GUI. The new features include settings that affect performance, such as toggling between TCP and UDP NFS versions 2 and 3. Other Server for NFS server settings include those that affect how file names are presented to NFS clients, such as allowing hidden files and allowing case sensitive lookups.



NOTE:

The NFS Server service needs to be restarted when changing these settings. Notify users when stopping and restarting the NFS Server service.

Use Remote Desktop to access the Management Console. Click **File Sharing, Microsoft Services for Network File System**. Click **Server for NFS**, then **Server Settings**.




Figure 67 Server for NFS screen, Server Settings tab

Installing NFS Authentication software on the domain controllers and Active Directory domain controllers

In mixed environments (for example, Windows 2000 and Windows NT 4.0), the NFS Authentication software must be installed on all Primary Domain Controllers (PDCs) and backup domain controllers (BDCs) that have Windows users mapped to UNIX users. This includes Active Directory domains. For instructions on setting up user mappings, see "NFS User and Group Mappings."



NOTE:

If the authentication software is not installed on all domain controllers that have user name mappings, including Primary Domain Controllers, Backup Domain Controllers, and Active Directory Domains, then domain user name mappings will not work correctly.

SFU 3.5 is used for NFS Authentication. SFU 3.5 can be downloaded at no charge from the Microsoft web site:

<http://www.microsoft.com/windows/sfu/downloads/default.asp>

To install the Authentication software on the domain controllers:

1. From the SFU 3.5 files, locate the directory named *SFU35SEL_EN*.
2. On the domain controller where the Authentication software is being installed use Windows Explorer to:
 - a. Open the shared directory containing *setup.exe*.

- b. Double-click the file to open it. Windows Installer is opened.



NOTE:

If the domain controller used does not have Windows Installer installed, locate the file InstMSI.exe on the SFU 3.5 directory and run it. After this installation, the Windows Installer program starts when opening setup.exe.

3. In the Microsoft Windows Services for UNIX Setup Wizard dialog box, click **Next**.
4. In the User name box, enter your name. If the name of your organization does not appear in the Organization box, enter the name of your organization there.
5. Read the End User License Agreement carefully. If you accept the terms of the agreement, click **I accept the terms in the License Agreement**, and then click **Next** to continue installation. If you click **I do not accept the License Agreement** (Exit Setup), the installation procedure terminates.
6. Click Custom Installation, and then click **Next**.
7. In the Components pane, click the down arrow next to Windows Services for UNIX, and then click **Entire component will not be available**.
8. Click the plus sign (+) next to Authentication Tools.
9. In the Components pane, click the plus sign (+) next to Authentication Tools.
10. Click **Server for NFS Authentication**, click **Will be installed on local hard drive**, and then click **Next**.
11. Follow the remaining instructions in the Wizard.



NOTE:

NFS users can be authenticated using either Windows domain accounts or local accounts on the Windows server. Server for NFS Authentication must be installed on all domain controllers in the domain if NFS users will be authenticated using domain accounts. Server for NFS Authentication is always installed on the computer running Server for NFS.

Installing SFU 3.5 from CD

Microsoft Services for Unix 3.5 CD has been included with the storage server and is needed for the following procedure.

To install the Authentication software on the domain controllers (CD Method):

1. Insert the Microsoft Windows Services for UNIX compact disc into the CD-ROM drive of the domain controller.
2. On the domain controller where the Authentication software is being installed, use Windows Explorer to:
 - a. Open the shared directory containing *OEMsetup.msi*.

b. Double-click the file to open it.
Windows Installer is opened.



NOTE:

If the domain controller used does not have Windows Installer available, locate the file InstMSI.exe on the SFU 3.5 CD and run it. After this installation, the Windows Installer program starts when opening setup.exe.

3. In the Microsoft Windows Services for UNIX Setup Wizard dialog box, click **Next**.
4. In the User name box, type your name. If the name of your organization does not appear in the Organization box, type the name of your organization there.
5. Read the End User License Agreement carefully. If you accept the terms of the agreement, click **I accept the terms in the License Agreement**, and then click **Next** to continue installation. If you click **I do not accept the License Agreement** (Exit Setup), the installation procedure terminates.
6. Click **Custom Installation**, and then click **Next**.
7. In the Components pane, click the down arrow next to Windows Services for UNIX, and then click **Entire component will not be available**.
8. Click the plus sign (+) next to Authentication Tools for NFS.
9. In the Components pane, click the plus sign (+) next to Authentication Tools.
10. Click **Server for NFS Authentication**, click **Will be installed on local hard drive**, and then click **Next**.
11. Follow the remaining instructions in the Wizard.



NOTE:

NFS users can be authenticated using either Windows domain accounts or local accounts on the Windows server. Server for NFS Authentication must be installed on all domain controllers in the domain if NFS users will be authenticated using domain accounts. Server for NFS Authentication is always installed on the storage server, which also runs the Server for NFS.

Understanding NTFS and UNIX permissions

When creating a NFS export, make sure that the NTFS permissions on the share allow the correct permissions that you want assigned to users/groups. The following helps clarify the translation between UNIX and NTFS permissions:

- The UNIX read bit is represented within NTFS as the List Folder/Read Data permission
- The UNIX write bit is represented within NTFS as the Create File/Write Data, Create Folders/Append Data, Write Attributes, and Delete Subfolders and Files permissions
- The UNIX execute bit is represented within NTFS as the Traverse Folder/Execute File permission

NFS file shares

NFS file shares are created in the same manner as other file shares, however there are some unique settings. Procedures for creating and managing NFS file shares are documented in the same sections as creating file shares for other protocols. See the “Folder and Share Management” chapter for more information.



NOTE:

NFS specific information is extracted from the “Folder and Share Management” chapter and duplicated below.

Complete share management is performed through the **Shares** menu option of the WebUI. Tasks include:

- Creating a new share
- Deleting a share
- Modifying share properties

Each of these tasks is discussed in this section.

Creating a new share

To create a new NFS file share:

1. From the WebUI main menu, select the **Shares** tab and then select the **Shares** option. The **Shares** page is displayed. From the **Shares** page, click **New**. The **General** tab of the **Create a New Share** page is displayed.




Figure 68 Create a New Share page, General tab

2. On the **General** tab, enter the share name and path. Select the **Unix (NFS)** client protocol check box.



NOTE:

Clear the Microsoft SMB option if you do not want to allow SMB access to the share.



NOTE:

NFS service does not support the use of spaces in the names for NFS file shares. NFS translates any spaces in an export into an underscore character. If you plan to use the same name when sharing a folder through SMB, and then exporting it through NFS, do not put spaces in the SMB share name.

To create a folder for the share, select the indicated box and the system will create the folder at the same time it creates the share.

3. Click the **NFS Sharing** tab to enter NFS specific information. See "Modifying Share Properties" for information on this tab.
4. After all share information is entered, click **OK**.

The default NFS share properties are **All Machines read only with root and anonymous access disallowed**. See the section, "Modifying Share Properties" in this chapter to change the default permissions.

Deleting a share



CAUTION:

Before deleting a share, warn all users to exit that share. Then confirm that no one is using the share.

To delete a share:

1. From the **Shares** page, select the share to be deleted, and then click **Delete**.
2. Verify that this is the correct share, and then click **OK**.

Modifying share properties

To change share settings:

1. From the **Shares** page, select the share to modify, and then click **Properties**.

The **General** tab of the **Share Properties** page is displayed.




Figure 69 Share Properties page, General tab

The name and path of the selected share is displayed.

2. To enter or change client protocol information, select the **UNIX (NFS)** client type box, and then click the **UNIX Sharing** tab.




Figure 70 UNIX Sharing tab

3. From the **UNIX Sharing** tab of the **Share Properties** page,
 - a. Indicate the allowed clients.
Select the machine to include in the **Select a group** box or manually enter the NFS client computer name or IP address. Then click **Add**.
 - b. Indicate the access permissions.
Select the machine from the main user display box and then select the appropriate access method from the **Access permissions** drop down box.
The types of access are:
 - **Read-only**—Use this permission to restrict write access to the share.
 - **Read-write**—Use this permission to allow clients to read or write to the share.
 - **No access**—Use this permission to restrict all access to the share.
4. Select whether or not to allow root access. Select the **Allow root access** checkbox to add the root permission.
 - **Read-only + Root**—Use this permission to restrict write access to the share. Use this permission to assign administrative access to the share. This allows the client computer to have root access to the NFS share. Map the UNIX root user to the Windows user Administrator. Also, map the group to which this UNIX root belongs to, to the Windows group Administrator.
 - **Read-write + Root**—Use this permission to allow clients to read or write to the share. Use this permission to assign administrative access to the share. This will allow the client computer to have root access to the NFS share. Map the UNIX root user to the Windows user Administrator. Also, map the group that this UNIX root belongs to, to the Windows group Administrator.
5. After all UNIX sharing information is entered, click **OK**.

Anonymous access to an NFS share

It may be desirable to add anonymous access to a share. An instance would be when it is not desirable or possible to create and map a UNIX account for every Windows user. A UNIX user whose account is not mapped to a Windows account is treated by Server for NFS as an anonymous user. By default, the user identifier (UID) and group identifier (GID) is -2.

For example, if files are created on an NFS Share by UNIX users whose are not mapped to Windows users, the owner of those files are listed as anonymous user and anonymous group, (-2,-2).

By default, Server for NFS does not allow anonymous users to access a shared directory. When an NFS share is created, the anonymous access option can be added to the NFS share. The values can be changed from the default anonymous UID and GID values to the UID and GID of any valid UNIX user and group accounts.

When allowing anonymous access to an NFS Share, the following must be performed by a user with administrative privileges due to Windows Storage Server 2003 security with anonymous users and the Everyone group.

1. From the WebUI, select **Maintenance**.
2. Click **Remote Desktop**. Log on to the storage server.
3. Click **Start >Control Panel > Administrative Tools**, and then click **Local Security Policy**.
4. In Security Settings, double-click Local Policies, and then click Security Options.
5. Right-click **Network access: Let Everyone permissions apply to anonymous users**, and then click **Properties**.
6. To allow permissions applied to the Everyone group to apply to anonymous users, click **Enabled**. The default is **Disabled**.
7. Restart the NFS server service. From a command prompt, enter net stop nfssvc. Then enter net start nfssvc. Notify users before restarting the NFS service.
8. Assign the Everyone group the appropriate permissions on the NFS Share.
9. Enable anonymous access to the share.

To enable anonymous access to an NFS share, do the following.

1. Open Windows Explorer by clicking **Start > Run**, and entering explorer.
2. Navigate to the NFS share.
3. Right-click the NFS Share, and then click **Properties**.
4. Click **NFS Sharing**.
5. Select the the **Allow Anonymous Access** checkbox.
6. Change from the default of -2,-2, if desired.
7. Click **Apply**.
8. Click **OK**.

Encoding Types

Encoding types can be selected using the WebUI. These include the default ANSI as well as EUC-JP. Other encoding types include:

- ANSI (default)
- BIG5 (Chinese)
- EUC-JP (Japanese)
- EUC-KR (Korean)
- EUC-TW (Chinese)
- GB2312-80 (Simplified Chinese)
- KSC5601 (Korean)
- SHIFT-JIS (Japanese)

If the option is set to ANSI on systems configured for non-English locales, the encoding scheme is set to the default encoding scheme for the locale. The following are the default encoding schemes for the indicated locales:

- Japanese: SHIFT-JIS
- Korean: KS C 5601-1987
- Simplified Chinese: GB
- Traditional Chinese: BIG5

NFS only

Microsoft Services for NFS allows the option of setting up NFS Shares for NFS access only.

The NFS Only option provides faster NFS performance and is intended for NFS clients only. The executable file, *nfsonly.exe*, allows a share to be modified to do more aggressive caching to improve NFS performance. This option can be set on a share-by-share basis. Do not use this function on any file share that can be accessed by any means other than by NFS clients, as data corruption can occur.

The syntax of this command is:`nfsonly <sharename> [/enable|/disable]`

- Sharename is the name of the NFS share
- The /enable option turns on NfsOnly for the specified share
- The /disable option turns off NfsOnly for the specified share

Restart the NFS service after setting up an NFS Only share. Notify users when the NFS service is restarted.

NFS protocol properties settings

Enter and maintain parameter settings for the NFS protocol through the WebUI in the **NFS Properties** page. To access the **NFS Properties** page, click **Shares**, **Sharing Protocols**. Then, select the **NFS Protocol** radio button and click **Properties**.

The **NFS Properties** page is displayed.

The screenshot shows the 'Shares' tab selected in the top navigation bar. Below it, the 'Sharing Protocols' section is active. The 'NFS' tab is selected. Under 'NFS', there are four main items: 'NFS Async/Sync Settings' (with a description of toggling between asynchronous and synchronous writes), 'NFS Locks' (with a description of managing service locks), 'Client Groups' (with a description of managing NFS Client Groups), and 'User and Group Mappings' (with a description of managing user name mapping). At the bottom left is a 'Back' button.

Figure 71 NFS Sharing Protocols page

NFS properties include:

- Async/Sync Settings (not available on all models)
- NFS Locks
- Client Groups
- User and Group Mappings

Settings for asynchronous/synchronous writes and service locks are discussed together in the following paragraphs of this chapter.

Client groups and user and group mappings are each discussed in separate sections later in this chapter.

NFS async/sync settings



NOTE:

The NFS async/sync settings feature is not available on models.

As mentioned in a previous section, there are two versions of NFS: Version 2 and Version 3. Version 3 supports additional file operations that Version 2 did not have, such as asynchronous file operations.

To indicate whether to use asynchronous or synchronous write settings:

1. From the WebUI, access the **NFS Protocol Properties** page by clicking **Shares**, **Sharing Protocols**. Select **NFS Protocol**, and then click **Properties**.

The **NFS Properties** page is displayed.

2. On the **NFS Properties** page, select **NFS Async/Sync Settings**.

The **NFS Async/Sync Settings** page is displayed.

3. Select the desired write setting. The default setting is Synchronous writes.



NOTE:

Using synchronous writes allows for greater data integrity. Asynchronous writes will increase performance but will reduce data integrity as the data is cached before being written to disk. Changing the write state causes the NFS service to be restarted. Notify users before toggling this setting.




Figure 72 NFS Async/Sync Settings page

NFS locks

NFS supports the ability to lock files. File locking helps prevent two or more users from working with the same files at the same time.

NFS locking depends on the software application components to manage the locks. If an application does not lock a file or if a second application does not check for locks before writing to the file, nothing prevents the users from overwriting files.

To enter locking parameters:

1. From the WebUI, access the **NFS Protocol Properties** page by clicking **Shares**, **Sharing Protocols**. Select **NFS Protocol**, and then click **Properties**.

The **NFS Properties** menu is displayed.

2. In the **NFS Properties** page, click **NFS Locks**.

The **NFS Locks** page is displayed. (See [Figure 73](#)).

All clients that have locks on system files are listed in the **Clients that hold locks** box.

3. To manually clear locked files, select the client from the displayed list, and then click **OK**.
4. To indicate the amount of time after a system failure that locks are kept active, enter the number of seconds in the **Wait period** box.

The storage server keeps the locks active for the specified number of seconds, while querying the client to see if it wants to keep the lock. If the client responds within this time frame, the lock is kept active. Otherwise, the lock is cleared.

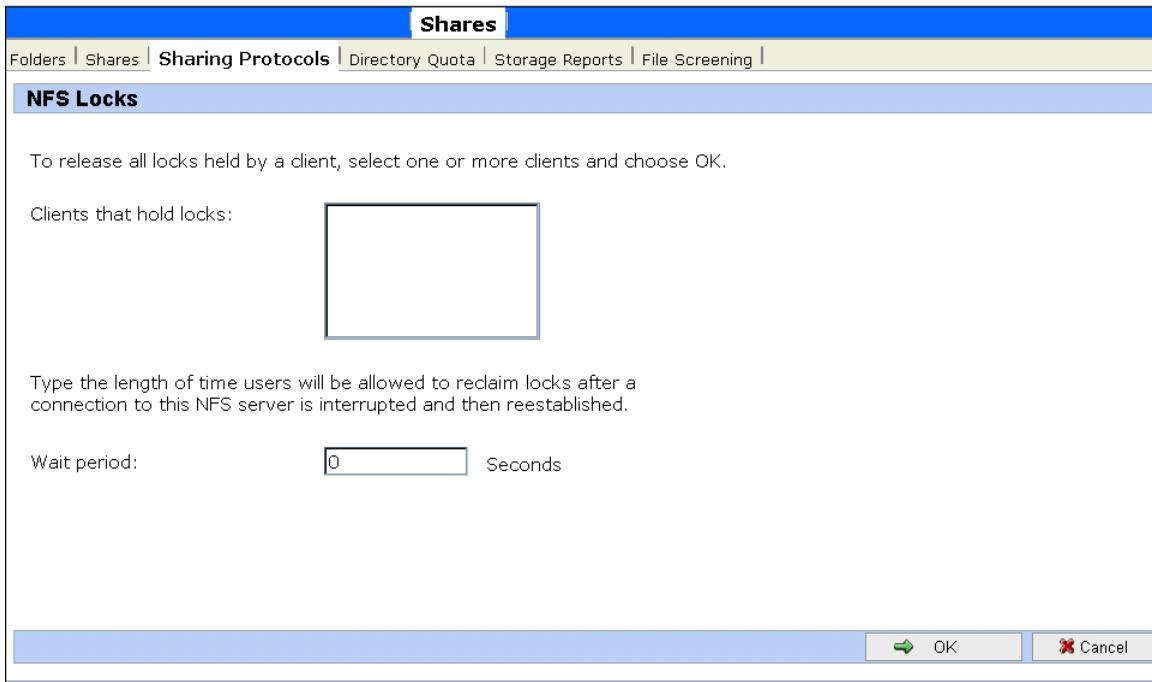


Figure 73 NFS Locks page

NFS client groups

The Client Groups feature gives administrators a method of assigning access permissions to a set of clients. The administrator creates a client group, gives it a name, and then inserts clients into the group by client name or IP address. After the client group is created, the administrator adds or removes permissions for the entire group, instead of allowing or denying access for each individual client machine.

Proper planning includes control over the naming conventions of client groups and users. If the client group is given the same name as a client, the client is obscured from the view of the server. For example, assume that a client d4 exists. If a client group called d4 is created, permissions can no longer be assigned to just the client d4. Any reference to d4 now refers to client group d4.

To manage NFS client groups:

- From the WebUI, access the **NFS Protocol Properties** page by clicking **Shares, Sharing Protocols**. Click **Client Groups**.

The **NFS Client Groups** page is displayed.

The screenshot shows a software interface titled "Shares" at the top. Below it is a navigation bar with links: Folders, Shares (which is the active tab), Sharing Protocols, Directory Quota, Storage Reports, and File Screening. The main content area is titled "NFS Client Groups". A message says, "Select a NFS Client Group, then choose a task. To create a new Client Group, choose New...". On the left, there's a list titled "NFS Client Groups" with one item: "ClientGroup1". On the right, a vertical column of buttons labeled "Tasks" includes "New...", "Delete", and "Edit...". At the bottom left is a "Back" button.

Figure 74 NFS Client Groups page

The following tasks are available:

- Adding a new client group
- Deleting a client group
- Editing client group information

Adding a new client group

To add a new client group:

1. From the **NFS Client Groups** page, click **New**.

The **New NFS Client Group** page is displayed.




Figure 75 New NFS Client Group page

2. Enter the name of the new group.
3. Enter the client name or their IP address.
4. Click **Add**. The system adds the client to the displayed list of members.
5. To remove a client from the group, select the client in the **Members** box, and then click **Remove**.
6. After all clients have been added to the group, click **OK**.

Deleting a client group

To delete a group:

1. From the **NFS Client Groups** page, select the group to delete and click **Delete**.
A verification screen is displayed.
2. Confirm that this is the correct group, and then click **OK**.

Editing client group information

To modify the members of an existing client group:

1. From the **NFS Client Groups** page, select the group to modify, and then click **Edit**.
The **Edit NFS Client Group** page is displayed. Current members of the group are listed in the **Members** box.




Figure 76 Edit NFS Client Groups page

2. To add a client to the group, enter the client name or IP address in the **Client name** box, and then click **Add**. The client is automatically added to the **Members** list.
3. To delete a client from the group, select the client in the **Members** list, and then click **Remove**. The client is removed from the list.
4. After all additions and deletions are completed, click **OK**.

NFS user and group mappings

When a fileserver exports files within a homogeneous environment, there are no problems with authentication. It is a simple matter of making a direct comparison to determine whether the user should be allowed access to the file, and what level of access to allow.

However, when a fileserver works in a heterogeneous environment, some method of translating user access is required. User mapping is the process of translating the user security rights from one environment to another.

User name mapping is the process of taking user and group identification from one environment and translating it into user identification in another environment. In the context of UNIX and NFS, user and group identification is a combination of a user ID (UID) and group ID (GID). In Windows environments, user identification is a Security ID (SID) or, in Windows Storage Server 2003, a Globally Unique Identifier (GUID).

The server grants or denies access to the export based on machine name or IP address. However, after the client machine has access to the export, user-level permissions are used to grant or deny access to user files and directories.

The storage server is capable of operating in a heterogeneous environment, meaning that it is able to work with both UNIX and Windows clients. Because the files are stored in the native Windows NT file system, the server has to map the UNIX users to Windows users to determine the user access level of the files.



NOTE:

User mapping is not designed to address existing user database problems in the existing environment. All UIDs and GIDs must be unique across all NIS (Network Information Service) domains and all user names must be unique across all Windows NT domains.

The storage server supports mappings between one or more Windows domains and one or more NIS domains. The mappings can be set up in the Services for NFS management console by clicking the simple mappings **Add** button, and then selecting the Windows domain names and the corresponding NIS domains.

Types of mappings

There are three types of mappings. These mappings are listed below in order of the most complex (with the greatest level of security) to the least complex (easiest to manage, but with little security):

- Explicit mappings
- Simple mappings
- Squashed mappings

Explicit mappings

Explicit mappings are created by the administrator to link Windows and UNIX users. They override simple mappings and are used to map users on the different systems that have unique names.

Simple mappings

Simple mapping is a direct comparison of user names on the Windows system and the UNIX system. If the names match, the user is assumed to be authentic, and appropriate share access is granted. Simple mapping is an option that the administrator must turn on if it is to be used.

Squashed mappings

If the NFS server does not have a corresponding UID or GID or if the administrator has set other conditions to filter out the user, a process called squashing takes effect. Squashing is the conversion of an unmapped or filtered user to an anonymous user. This anonymous user has very restricted permissions on the system. Squashing helps administrators manage access to their exports by allowing them to restrict access to certain individuals or groups and to squash all others down to restricted (or no) access. Squashing enables the administrator to allow permissions instead of denying access to all the individuals who are not supposed to have access.

Figure 77 is a diagram showing an example of how the mapping server works for an ls -al command.




Figure 77 Mapping server “ls -al” command example

A double translation, as illustrated in Figure 77, is sometimes necessary because some commands return user ID information. For example, if the NFS request issued was an ls -al command, the return listing of files contains user information (the user and group that own the file). The ls -al command is a UNIX command. It returns a long or full listing of all files. Because this information is contained in a Windows NT Access Control List (ACL), it is not UNIX ready. The ACL information has to be converted back to UNIX UIDs and GIDs for the UNIX systems to understand and display the user information.

This second translation is not done for commands that do not return user information. For example, if the NFS request were just to read data from or write data to a file, the second translation would not be performed because there is no returning user information.

User name mapping best practices

Below is a brief list of suggested practices:

- **Back up user and group mappings**

To avoid loss of complex advanced mappings in the case of a system failure, back up the mappings whenever the mappings have been edited or new mappings have been added.

- **Map consistently**

Groups that are mapped to each other should contain the same users, and members of the groups should be properly mapped to each other to ensure proper file access.

Example using User1 and Group1:

- Verify that the Windows User1 is mapped to the corresponding UNIX User1.
- Verify that the Windows Group1 is mapped to the corresponding UNIX Group1.
- Verify that User1 is a member of Group1 on both Windows and UNIX.

- **Map properly**

- PCDATA cannot be inserted: Valid UNIX users should be mapped to valid Windows users.
- PCDATA cannot be inserted: Valid UNIX groups should be mapped to valid Windows groups.
- PCDATA cannot be inserted: The mapped Windows user must have the “Access this computer from the Network privilege,” or the mapping will be squashed.
- PCDATA cannot be inserted: The mapped Windows user must have an active password, or the mapping will be squashed.

Creating and managing user and group mappings



NOTE:

The following sections are for a stand alone configuration. In a clustered environment, clicking User and Group Mappings displays a log in screen for the Services for NFS Administrator.

To set up and manage user name mappings:

1. From the WebUI, click **Shares, Sharing Protocols**. Click **NFS Protocol**, and then click **Properties**.

The **NFS Properties** page is displayed.

2. In the **NFS Properties** page, click **User and Group Mappings**.

There are four tabs in the **User and Group Mappings** page:

- **General information**—Sets the mapping information source, which is either NIS or password and group files.
- **Simple Mapping**—Indicates whether simple mappings are being used.
- **Explicit User Mapping**—Lists exceptional user mappings that will override the simple user mappings.
- **Explicit Group Mapping**—Lists exceptional group mappings that will override the simple group mappings.

Each of these tabs is discussed in the following sections.

3. Enter mapping information on the appropriate tabs, and then click **OK**.

General tab

The user name mapping server translates the UNIX users into Windows users so that the server can determine user access rights to the data.

Within this initial page, indicate whether the source of mapping information is an NIS server, or is a special file with password and group information.




Figure 78 User and Group Mappings page, General tab

From the **General** tab of the **User and Group Mappings** page:

1. If an NIS server is being used:
 - a. Select **Use NIS server**.
 - b. Enter the NIS domain name.
 - c. Enter the NIS server name. This field is optional, but recommended. In the **Hours** and **Minutes** fields, indicate how often the system will connect to the NIS domain to update the user list.
2. If custom password and group files are being used:
 - a. Select **User password and group files**.
 - b. Enter the path and name of the password file.
 - c. Enter the path and name of the group file.
3. After this basic information is entered, click **OK**.

Simple mapping tab

Simple (or implicit) mapping is the first level of user name mapping. In simple mode, user and group names that match exactly in name are automatically equated.

While simple mappings are the most easily managed and are the most forthright type of map, security problems can arise. For example, if a UNIX user is coincidentally an exact match of a Windows user, the system will equate them and an inadvertent mapping will occur, granting a user inappropriate access.

- To use simple mappings, the feature must be enabled. If this feature is disabled, the administrator must manually create an explicit map for each user.
- To enable simple mapping, select the **Enable Simple Mapping** option, and then select the Windows domain name.




Figure 79 User and Group Mappings page, Simple Mapping tab

Explicit user mapping tab

Explicit (or advanced) mappings allow the administrator to map any user or group manually to any other user and group. Advanced mappings override simple mappings, giving administrators the capability of using simple mapping for most users and then using advanced mappings for the users with unique names on the different systems. Alternatively, simple mapping can be disabled completely, relying solely on explicit mappings. Explicit mappings create the most secure mapping environment.

Security issues seen in simple mappings do not exist in explicit mappings. Explicit user mappings specifically correlate two users together, thus preventing the inadvertent mapping.

To enter explicit user mappings, click the **Explicit User Mapping** tab. (See [Figure 80](#)).




Figure 80 User and Group Mappings page, Explicit User Mapping tab

To create explicit user mappings:

1. Click the **List UNIX Users** button to populate the UNIX users box.
2. To map a local Windows user to a UNIX user, highlight the **Windows user** in the Windows local users box and highlight the UNIX user that you want to map, and then click **Add**. The **Explicitly mapped users** box at the bottom of the screen is populated with the new mappings. Repeat this process until all desired users have been mapped.
3. To map a domain Windows user to a UNIX user, enter the domain and the user name in the box in the middle of the screen (use the Domain\username format) and highlight the UNIX user that you want to map, and then click **Add**. The map is added to the **Explicitly mapped users** box at the bottom of the page. Repeat this process until all desired users have been mapped.
4. To map multiple Windows users to one UNIX user, one of the mapped Windows users must be set as the primary mapping. To indicate which user map is the primary mapping, highlight the desired map in the **Explicitly mapped users** box, and then click the **Set Primary** button.
5. To delete a map, highlight the map in the **Explicitly mapped users** box, and then click the **Remove** button.
6. After all entries are completed, click **OK** to activate the new entries.

Explicit group mapping tab

To enter explicit group mappings, select the **Explicit Group Mapping** tab. [Figure 81](#) is an example of the **Explicit Group Mapping** tab.

Explicit mappings allow the administrator to map any user or group manually to any other user and group. Explicit mappings override simple mappings, giving administrators the capability of using simple

mapping for most groups and then using explicit mappings to make changes to simple mappings. Simple mapping can be turned off for greater security.




Figure 81 User and Group Mappings page, Explicit Group Mapping tab

To create explicit group mappings:

1. Click the **List UNIX Groups** button to populate the **UNIX Groups** box.
2. To map a local Windows group to a UNIX group, highlight the Windows group in the Windows local groups box and highlight the UNIX group to map, and then click **Add**. The **Explicitly mapped groups** box at the bottom of the screen is populated with the new mappings. Repeat this process until all desired groups have been mapped.
3. To map a domain Windows group to a UNIX group, enter the domain and the group name in the box in the middle of the screen (use the Domain\groupname format) and highlight the UNIX group to map, and then click **Add**. The map is added to the **Explicitly mapped groups** box at the bottom of the page. Repeat this process until all desired groups have been mapped.
4. To map multiple Windows groups to one UNIX group, one of the Windows groups must be set as the primary mapping. Therefore, to indicate which group map is the primary mapping, highlight the desired map in the **Explicitly mapped groups** box, and then click the **Set Primary** button.
5. To delete a map, highlight the map in the **Explicitly mapped groups** box, and then click **Remove**.
6. After all entries are completed, click **OK** to activate the new entries.

Backing up and restoring mappings

The user name-mapping server has the capability to save and retrieve mappings from files. This capability is useful for backing up mapping settings prior to making changes and for exporting the mapping file from one server to others, using the same mapping information.

The user name-mapping server can save existing mappings to a file, or load them from a file and populate the mapping server. This feature is found in the **Map Maintenance** tab of the **User Name Mapping** screen, as shown in [Figure 82](#).

Use **Remote Desktop** to access the **Management Console**, click **File Sharing, Microsoft Services for Network File System**. Click **User Name Mapping**, then **Map Maintenance**.




Figure 82 User Name Mapping screen, Map Maintenance tab

Backing up user mappings

1. Click the **Map Maintenance** tab on the **User Name Mapping** screen.
2. Type the path and name of the file to be used for backup in the File path and name field or click **Browse** to locate the file.



NOTE:

If the file is being created for the first time, follow these steps:

1. Browse to the target directory.
2. Right-click in the file listing pane, select New, Text Document. Enter a name for the file, and then press Enter.
3. Double-click the new file to select it.
4. Click Backup.

Restoring user mappings

User mappings can be restored using the following procedures.

1. Click the **Map Maintenance** tab on the **User Name Mapping** screen.
2. Type the path and name of the file in the File path and name field, or click **Browse** to locate the file.
3. After locating the file, click **Restore**.

Creating a sample NFS file share

HP recommends performing the following tests to verify that the setup of the shares, user mappings, and permissions grant the desired access to the NFS shares.

1. Create an NFS share. NFS Shares are All Machines, read-only by default.
See "NFS File Shares" earlier in this chapter for information on creating shares.
2. Create NFS client groups if desired. See "NFS Client Groups" earlier in this chapter.
3. Verify that the NFS share exists.

Use Remote Desktop to log into the storage server and access the command line interface:

`nfsshare <sharename>` (sharename represents the name of the share)

4. Map a user. When creating Active Directory/Domain mappings, ensure that the NFS Authentication software is installed on the domain controllers that have user name mappings. See "Installing NFS Authentication Software on the Domain Controllers and Active Directory Domain Controllers" section. Also, see "User and Group Mappings" in this chapter for instructions on setting up user name mappings. When planning to allow only anonymous access to an NFS share, setting up user name mappings is not necessary.

5. Verify the NTFS permissions are correct on the NFS share. If the NFS share was assigned All Machines read write, the NTFS ACLs on the NFS share must allow read/write permissions for the user or group. Example: `f:\share1` is the name of the NFS share and `share1` has All Machines read write permissions. Verify that the NTFS permissions on `f:\share1` are List Folder/Read Data, Create File/Write Data, Create Folders/Append Data, Write Attributes, and Delete Subfolders and Files. This can be verified by opening up Windows Explorer on the storage server desktop and right-clicking `f:\share1`, and then clicking **Properties**. Next, click the **Security** tab. Then click **Advanced**. Highlight the user or group that permissions are being assigned to, and then click **Edit**. There are checkboxes next to the NTFS permissions that are assigned. Make sure mapped users and groups correlate to the users or groups that have the NTFS permissions assigned. See the section "Understanding NTFS and UNIX Permissions" in this chapter for more information.

6. Verify that the mappings exist.

Use Remote Desktop to log in to the storage server and access the command line interface:

`mapadmin list -all`

7. On the Linux/UNIX system, use the mapped user to create a file.

- a. As the root user, mount the share:
`mount -t nfs <nfs server IP address:/nfs share> /mount point`
- b. Log in as a mapped user.
- c. Change directories to the mount-point directory.
- d. Create the file as the mapped user (example: `file1`).

8. Verify that the same permissions are set up for the user on both the UNIX side and the Windows side.

- a. List the permissions on the UNIX side:

```
ls -l /mount-point/file1
```

(Example screen display: -r-r--- unixuser1 unixgroup1)

- b. List the permissions on the Windows side: (change to the *nfs* share directory)

From a command line interface accessed from Remote Desktop on the storage server:

```
cacls file1
```

(Example display: DOMAIN1\Windowsuser1:R)

- c. Compare and verify the permissions from UNIX and Windows.

Remote Access

Using Remote Desktop

In addition to the WebUI, Remote Desktop is available for remote administration of Services for UNIX. This service let users connect to machines, log on, and obtain command prompts remotely. See [Table 14](#) for a list of commonly used commands.



CAUTION:

Two open sessions of Remote Desktop are allowed to operate at the same time. After completing an application do not use the window close feature (to close that session of Remote Desktop. Select Start > Log Off Administrator to exit Remote Desktop.

Microsoft Remote Desktop can be used to remotely access the storage server desktop. This provides the administrator flexibility to automate setups and other tasks. Services for NFS file-exporting tasks and other Services for NFS administrative tasks can be accomplished using Remote Desktop to access the Services for NFS user interface from the storage server desktop or from a command prompt.

Remote Desktop is included in the WebUI of the storage server. To open a Remote Desktop session, from the WebUI, click **Maintenance, Remote Desktop**. See the “Remote Access Methods and Monitoring” chapter for information on setting up and using Remote Desktop.

[Table 14](#) describes some common Services for NFS commands.

Table 14 Command Line Interface Command Prompts

Command	Function
nfssstat /?	Learn about viewing statistics by NFS operation type.
showmount /?	View the format of the command to display NFS export settings on NFS servers.
showmount -a	View users who are connected and what they currently have mounted.
showmount -e	View exports from the server and their export permissions.
rpcinfo /?	Learn how to display Remote Procedure Call (RPC) settings and statistics.
mapadmin /?	View how to add, delete, or change user name mappings.
nfsshare /?	Learn how to display, add, and remove exported shares.

Using Telnet Server

**NOTE:**

Telnet Server is not available on all models.

Telnet is a UNIX command line utility. The Telnet service is included on the storage server, but, by default, is not activated. To use Telnet services, see the information in the "Remote Access Methods and Monitoring" chapter.

**NOTE:**

The version of Telnet Server that can be administered by the SFU MMC is the standard Telnet Server found on Windows 2003.

Using Remote Shell Service

**NOTE:**

Remote Shell Service is not available on all models.

The Remote Shell is a UNIX method for allowing UNIX users to run commands remotely. It can be used in a fashion similar to Telnet, or can be used to directly invoke a remote command. Remote Shell service is not activated by default.

**NOTE:**

Remote Shell Service is not cluster-aware.

Interix

**NOTE:**

The information in the remainder of this chapter refers to storage servers using SFU 3.5.

Interix is a full application execution subsystem that allows administrators to compile and natively run UNIX programs and scripts on the storage server. It includes a full set of UNIX utilities and shells, support for a single-rooted file system, and a software development kit (SDK) for porting applications.

Shells

Both Korn and C shells are available in the Interix subsystem. Both shells behave as they do in a UNIX environment, making it much easier to port scripts from UNIX to Windows.

Programming Languages

The Interix environment includes support for Perl, C, fortran77, and C++. In addition, there are updated versions of the GNU programming languages and tools, optimized for SFU, as part of the GNU SDK.

Enabling setuid behavior for Interix programs

According to the POSIX standard, a file has permissions that include bits to set a UID (setuid) and set a GID (setgid) when the file is executed. If either or both bits are set on a file, and a process executes that file, the process gains the UID or GID of the file. When used carefully, this mechanism allows a nonprivileged user to execute programs that run with the higher privileges of the file's owner or group. When used incorrectly, however, this can present security risks by allowing nonprivileged users to perform actions that should only be performed by an administrator. For this reason, Windows Services for UNIX Setup does not enable support for this mechanism by default.

You should enable support for setuid behavior only if you are sure you will be running programs that require support for this behavior. By default, support for setuid is not available in Interix. To enable this behavior, search for "enable setuid mode bits" in the Windows Services for UNIX help and follow the instructions in the help topic.

8 NetWare File System Management

File and Print Services for NetWare (FPNW) is one part of the Microsoft software package called Services for NetWare. The most common use of the NetWare network operating system is as a file and print server. FPNW eases the addition of the storage server into a mixed infrastructure by providing a NetWare user interface (UI) to a Windows Storage Server 2003-based server; administrators and users see their same, familiar NetWare UI. Additionally, the same logon for clients is maintained without a need for any client configuration changes.

This service also provides the ability to create Novell volumes, which are actually NTFS shares, from which users can map drives and access resources. Novell Login scripts are supported on the storage server or through an existing NDS (Novell Directory Services) account.



NOTE:

FPNW is not a clusterable protocol. With FPNW on both nodes of a cluster, the shares do not fail over because the protocol is not cluster-aware.



NOTE:

IPX/SPX protocol is required on the Novell servers.

Installing Services for NetWare

The installation of FPNW on the storage server allows for a smooth integration with existing Novell servers. FPNW allows a Windows Storage Server 2003 based server to emulate a NetWare file and print server to users, clients, and administrators. This emulation allows authentication from Novell clients, the use of Novell logon scripts, the creation of Novell volumes (shares), the use of Novell file attributes, and many other Novell features.

Information on Microsoft Directory Synchronization Services and the File Migration Utility can be found at:

<http://www.microsoft.com/WINDOWS2003/guide/server/solutions/NetWare.asp>

To install Services for NetWare:

1. From the desktop of the storage server, select **Start > Settings > Network Connections > Local Area Connection**, and then right-click **Properties**.
2. Click **Install**.

The **Select Network Component Type** dialog box is displayed.




Figure 83 Local Area Connection Properties page, Install option

3. Click **Service**, and then click **Add**.
4. Click the **Have Disk** icon, and then navigate to the location of **Services for NetWare**.
Services for NetWare is located in the path: `c:\hpnas\components/SFN5.02/fpnw/netsfn.inf`.
5. Select the `NETSFNTSRV` file, and then click **OK**.

File and Print Services for NetWare should now be displayed as an option to install.

6. Select **File and Print Services for NetWare**, and then click **OK**.




Figure 84 Installing File and Print Services for NetWare

Managing File and Print Services for NetWare

To access FPNW:

- From the desktop of the storage server, select **Start > Settings > Control Panel > Administrative Tools > Server Manager**.
- Select **FPNW**, and then click **Properties**.




Figure 85 File and Print Services for NetWare dialog box

- Enter an FPNW Server Name and Description.

This server name must be different from the server name used by Windows or LAN Manager-based clients. If changing an existing name, the new name is not effective until stopping and restarting FPNW. For example, in Figure 85 the Windows server name is Alamo and the FPNW server name is Alamo_FPNW.

- Indicate a Home directory root path.

This path is relative to where the Sysvol volume is installed. This is the root location for the individual home directories. If the directory specified does not already exist, it must first be created.

- Click **Users** to:

See connected users, disconnect users, send broadcast messages to all users connected to the server, and to send a message to a specific user.

- Click **Volumes** to:

See users connected to specific volume and to disconnect users from a specific volume.

- Click **Files** to:

View open files and close open files.

Creating and managing NetWare users

To use Services for NetWare, the Novell clients must be entered as local users on the storage server.

Adding local NetWare users

1. From the storage server desktop, click the **Management Console** icon, click **Core Operating System**, and then click **Local Users and Groups**.
2. Right-click the **Users** folder, and then click **New User**.




Figure 86 New User dialog box

3. Enter the user information, including the user's User name, Full name, Description, and Password.
4. Click **Create**.
5. Repeat these steps until all NetWare users have been entered.

Enabling local NetWare user accounts

1. In the **Users** folder (MC, Core Operating System, Local Users and Groups), right-click an NCP client listed in the right pane of the screen, and then click **Properties**.
2. Click the **NetWare Services** tab.




Figure 87 NetWare Services tab

3. Select **Maintain NetWare compatible login**.
4. Set other NetWare options for the user, and then click **OK**.



NOTE:

The installation of File and Print Services for NetWare also creates a supervisor account, which is used to manage FPNW. The supervisor account is required if the storage server was added as a bindery object into NDS.

Managing NCP volumes (shares)

NCP file shares are created the same way as other file shares; however, there are some unique settings. NCP shares can be created and managed using Server Manager.



NOTE:

NCP shares can be created only after FPNW is installed. See the previous section “[Installing Services for NetWare](#)” for instructions on installing SFN.

Creating a new NCP share

To create a new file share:

1. From the storage server desktop, select **Start > Settings > Control Panel > Administrative Tools > Server Manager**.
2. Select **File and Print Service for NetWare > Shared Volumes**.
3. Click **Create Volume**.




Figure 88 Create Volume dialog box

4. Specify the volume name and path.
5. Click **Permissions** to set permissions.




Figure 89 Access Through Share Permissions dialog box

6. Click **Add** to add additional users and groups, and to set their permissions.




Figure 90 Add Users and Groups dialog box

7. Highlight the desired user or group, and then click **Add**.
8. Select the Type of Access in the drop down list.

Type of Access can also be set from the Access Through Share Permissions dialog box.

9. Click **OK** when all users and groups have been added.
10. Click **OK** in the **Create Volume** dialog box.
11. Click **Close**.

Modifying NCP share properties

To modify a file share:

1. From the storage server desktop, select **Start > Settings > Control Panel > Administrative Tools > Server Manager**.
2. Select **File and Print Services for NetWare > Shared Volumes**.
3. Highlight the volume to modify.
4. Click **Properties**.

9 Remote Access Methods and Monitoring

The HP ProLiant storage server ships with full remote manageability. Several methods of remote access are provided.

Web-based user interface

The storage server includes a web-based user interface (WebUI) for the administrator to remotely manage the machine. Of all of the remote access methods, the WebUI is the most intuitive and easiest to learn and use.

The WebUI permits complete system management, including system configuration, user and group management, shares management, UNIX file system management, and storage management.

To access the WebUI:

1. Launch a web browser.
2. In the URL field, enter:

`https://<your server machine name or IP address>:3202/`

Extensive procedural online help is included in the WebUI.

Remote Desktop

The storage server supports Remote Desktop, with a license for two concurrently running open sessions. Remote Desktop provides the same capabilities as being physically present at the server console.

Use Remote Desktop to access:

- The storage server desktop
- The Management Console
- A command line interface
- Backup software
- Antivirus programs
- Telnet Server

To access Remote Desktop from the WebUI, click **Maintenance, Remote Desktop**. For additional procedural information on Remote Desktop, see the “[Basic Administrative Procedures and Setup Completion](#)” chapter.

Telnet Server

Telnet Server is a utility that lets users connect to machines, log on, and obtain a command prompt remotely. Telnet Server is preinstalled on the storage server, but must be activated before use.



CAUTION:

For security reasons, the Telnet Server service must be restarted each time the server is restarted.

Enabling Telnet Server

Telnet Server can be enabled in two ways.

The first is to use Remote Desktop to access a command line interface, and then enter the following command:

```
net start tlntsvr
```

The Telnet Server service needs to be enabled prior to running this command. The service can be enabled by opening the services MMC:

1. Select **Start > Run**, and then enter services.msc.
2. Locate the Telnet service, right-click on it, and then select **Properties**.
3. In the Startup Type drop-down box, click **Manual**, and then click **OK**.

The second is to open the WebUI:

1. Click **Network**.
2. Click **Telnet**.
3. Select the **Enable Telnet access to this appliance** box.
4. Click **OK**.

Sessions information

The sessions screen provides the ability to view or terminate active sessions.

Integrated Lights-Out port



NOTE:

The Integrated Lights-out port capabilities are not supported on all storage server models. Refer to the QuickSpecs for a listing of specific model features.

The following information provides an overview of the Integrated Lights-Out port capabilities. For further information, refer to the *Integrated Lights-Out Port Installation and Users Guide* on the Documentation CD.

The Integrated Lights-Out port is an ASIC-based Web interface that provides remote management for the server.

Regardless of the state of the host operating system or the host CPU, complete capability for the server is available. The Integrated Lights-Out port is independent of the host server and its operating system. The Integrated Lights-Out port provides remote access, sends alerts, and performs other management functions, even when the host server operating system is not responding.

Features

The Integrated Lights-Out port provides the following features:



NOTE:

The remote client console must have a direct browser connection to the Integrated Lights-Out port without passing through a proxy server or firewall.

- Hardware based graphical remote console access
- Remote restart
- Server failure alerting
- Integration with Insight Manager
- Local Area Network (LAN) access through onboard NIC
- Browser support for Internet Explorer 5.50 or later
- Reset and failure sequence replay
- Auto configuration of IP address through domain name system (DNS) or Dynamic Host Configuration Protocol (DHCP)
- Virtual power button

Security features

- SSL encryption for login and network traffic
- User administration allows capability to define user profiles
- Event generation for invalid login attempts
- Logging of user action in the Event Log

Manage Users feature

The Manage Users feature allows those with supervisory access to add and delete users or to modify an existing user's configuration. Manage Users also lets the administrator modify:

- User name
- Logon name
- Password
- Simple network management protocol (SNMP) trap IP address
- Receive host OS generated SNMP traps
- Supervisor access
- Logon access
- Remote console access
- Remote server reset access

Manage Alerts feature

The Manage Alerts feature allows the user to:

- Select alert types received
- Generate a global test alert
- Generate an individual test alert
- Clear pending alerts
- Enable alerts

Refer to the *Integrated Lights-Out Port User Guide* for more information about the Integrated Lights-Out port features and functionality.

Integrated Lights-Out Port configuration

The Integrated Lights-Out port on the storage server is initially configured through the Rapid Startup Utility. SNMP is enabled and the Insight Management Agents are preinstalled.

The Integrated Lights-Out port comes with factory default settings, which the administrator can change. Administrators may want to add users, change SNMP trap destinations, or change networking settings. Refer to the *Integrated Lights-Out Port User Guide* for information about changing these settings.

There are several methods for performing Integrated Lights-Out port configuration changes:

- Web interface
- Integrated Lights-Out port configuration utility accessed by pressing **F8** during a system restart



NOTE:

You must connect locally with a monitor, keyboard, and mouse to utilize the F8 feature.

- Integrated Lights-Out port access using the default DNS name

Using the Integrated Lights-Out Port to Access the Storage Server

HP recommends using the web interface of a client machine to access the server remotely:

1. In the URL field of the web browser, enter the IP address of the Integrated Lights-Out port.



NOTE:

The iLO port can also be accessed from the HP Utilities tab of the WebUI by clicking the remote management link.

2. At the Integrated Lights-Out Account Login window, supply the username and password for the iLO and click **Login**.
3. Click the **Remote Console** tab.
The Remote Console Information screen is displayed.
4. Click the Remote Console choice in the menu on the left side of the screen.
5. Press **Ctrl-Alt-Del** to log into the console.
6. Supply an administrator username and password.

The storage server desktop is displayed.



NOTE:

The remote desktop feature of the iLO port requires a license key. The key is included with the product inside the Country Kit. See the iLO Advanced License Pack for activation instructions.

HP Insight Manager Version 7



NOTE:

The Integrated Lights-out port capabilities are not supported on all storage server models. Refer to the QuickSpecs for a listing of specific model features.

Some models of the storage server are equipped with the latest Insight Management Agents for Servers, allowing easy manageability of the server through HP System Management, HP OpenView, and Tivoli NetView.

Insight Manager is a comprehensive management tool that monitors and controls the operation of HP servers and clients. HP Insight Manager Version 7.0 or later is needed to successfully manage the storage server using the following components:

- Windows-based console application available on the Insight Manager 7 CD-ROM loaded on a separate client for storage servers
- Server or client based management data collection agents

Management agents monitor over 1,000 management parameters. Key subsystems make health, configuration, and performance data available to the agent software. The agents act upon that data by initiating alarms in the event of faults. The agents also provide updated management information, such as network interface or storage subsystem performance statistics.

10 Cluster Administration



NOTE:

Not all HP ProLiant Storage Servers can be clustered. Refer to the HP ProLiant Storage Server QuickSpecs to determine if your storage server can be clustered.

One important feature of the HP ProLiant Storage Server clusterable models is that they can operate as a single node or as a cluster. This chapter discusses cluster installation and cluster management issues.

Cluster overview

Two server nodes can be connected to each other and deployed as a no single point of failure (NSPOF) dual redundant cluster. The nodes are connected by a crossover cable and are each connected to network switches or hubs. This connection allows communication between the nodes to track the state of each cluster node. Each node sends out periodic messages to the other node; these messages are called heartbeats. If a node stops sending messages, the cluster service fails over any resources that the node owns to the other node. For example, if the node that owns the Quorum disk is shut down for any reason, its heartbeat stops. The other node detects the lack of the heartbeat and takes over ownership of the Quorum disk and the cluster.

Multi-node support beyond two nodes

Clusterable storage server devices may be deployed in multi-node clustering beyond two nodes. Refer to the associated Storage Array documentation to determine the number of nodes supported by the array under Windows Storage Server 2003. While the discussion presented in this guide addresses only two nodes, additional nodes may be added into the cluster. Considerations for additional fiber path connections and the private network should be made. In the case of the private network, a hub or switch is required since the cross over cable is no longer applicable.




Figure 91 Storage server cluster diagram

Cluster terms and components

Nodes

The most basic parts of a cluster are the servers, referred to as nodes. A server node is any individual computer in a cluster or a member of the cluster.

Resources

Hardware and software components that are managed by the cluster service are called cluster resources. Cluster resources have three defining characteristics:

- They can be brought online and taken offline.
- They can be managed in a cluster.
- They can be owned by only one node at a time.

Examples of cluster resources are IP addresses, network names, physical disk resources, and file shares.

Virtual servers

A virtual server is a cluster group that consists of a static IP Address resource and a Network Name resource. Several virtual servers can be created. By assigning ownership of the virtual servers to the different server nodes, the processing load on the storage servers can be distributed between the nodes of a cluster.

The creation of a virtual server allows resources dependant on the virtual server to fail over and fail back between the cluster nodes. File share and physical disks resources are assigned to the virtual server to ensure non disruptive service of file shares to the clients.

Failover

Failover of cluster groups and resources happens:

- when a node hosting the group becomes inactive. A shutdown of cluster service or a loss of power can cause a failover.
- when all of the resources within the group are dependent on one resource and that resource fails.
- when an administrator forces a failover.

A resource and all of its dependencies must be located in the same group so that if a resource fails over, all of its dependent resources fail over.

When a resource is failed over, the cluster service performs certain procedures. First, all of the resources are taken offline in an order defined by the resource dependencies. Secondly, the cluster service attempts to transfer the group to the next node on the preferred owners list. If the transfer is successful, the resources are brought online in accordance with the resource dependency structure.

The system failover policy defines how the cluster detects and responds to the failure of individual resources in the group. After a failover occurs and the cluster is brought back to its original state, failback can occur automatically based on the policy. After a previously failed node comes online, the cluster service can fail back the groups to the original host. The failback policy must be set before the failover occurs so that failback works as intended.

Quorum disk

Each cluster must have a shared disk called the Quorum disk. This physical disk in the common cluster disk array plays a critical role in cluster operations. The Quorum disk offers a means of persistent storage. The disk must provide physical storage that can be accessed by any node in the cluster. If a node has control of the quorum resource upon startup, it can initiate the cluster. In addition, if the node can communicate with the node that owns the quorum resource, it can join or remain in the cluster.

The Quorum disk maintains data integrity by:

- storing the most current version of the cluster database.
- guaranteeing that only one set of active communicating nodes is allowed to operate as a cluster.

Cluster concepts

[Figure 92](#) illustrates a typical cluster configuration with the corresponding storage elements. The diagram progresses from the physical disks to the file shares, showing the relationship between both the cluster elements and the physical devices underlying them. While the diagram only illustrates two nodes, the same concepts apply for multi node deployments.

Sequence of events for cluster resources

The sequence of events in the diagram includes:

1. Physical disks are combined into RAID arrays and LUNs.
2. LUNS are designated as basic disks, formatted, and assigned a drive letter via Disk Manager
3. Physical Disk resources are created for each basic disk inside Cluster Administrator.

4. Directories and folders are created on assigned drives.
5. Cluster components (virtual servers, file shares) are created, organized in groups, and placed within the folders using Cluster Administrator exclusively.




Figure 92 Cluster concepts diagram

Hierarchy of cluster resource components

The cluster components are referred to as resources and are placed together in groups. Groups are the basic unit of failover between nodes. Resources do not fail over individually, rather they fail over with the group in which they are contained.

In Figure 92 it is depicted as follows:

- Physical Disk resources are placed in a group and relate to the basic disk created through the WebUI. When a Physical Disk resource is created through Cluster Administrator, a corresponding group should be created for the resource to reside in. Groups are the basic unit of failover on a cluster.
- File share resources are placed in a group and relate to the actual directory on the drive on which the share is being created.
- An IP Address resource is formed in the group and relates to the IP address by which the group's virtual server is identified on the network.
- A Network Name resource is formed in the group and relates to the name published on the network by which the group is identified.
- A Virtual Server is a group containing an IP Address resource and a Network Name resource. File share and disk resources assigned to this virtual server group can transition from one node to the other during failover conditions.
- The Group is owned by one of the nodes of the cluster, but may transition to the other nodes during failover conditions.

The diagram illustrates a cluster containing two nodes. Each node has ownership of one group. Contained within each group are singular file shares that are known on the network by the associated Network Name and IP address. In the specific case of Node1, file share Eng1 relates to E:\Eng1. This file share is known on the network as \\Fileserver1\Eng1 with an IP address of 172.18.1.99. E:\Eng1 relates to the actual Basic Disk E: containing a directory Eng1.

For cluster resources to function properly, two very important requirements should be adhered to:

- Dependencies between resources of a group must be established. Dependencies determine the order of startup when a group comes online. In the above case, the following order should be maintained:

- File Share—Dependent on Physical Disk Resource
- NFS File Share—Dependent on Physical Disk Resource and Network Name
- Network Name—Dependent on IP Address

Failure to indicate the dependencies of a resource properly may result in the file share attempting to come online prior to the physical disk resource being available, resulting in a failed file share.

- Groups should have a Network Name resource and an IP Address resource. These resources are used by the network to give each group a virtual name. Without this virtual reference to the group, the only way to address a share that is created as a clustered resource is by node name. Physical node names do not transition during a failover, whereas virtual names do.

For example, if from a client a network share map F: was established and assigned to \\Node1\Eng1 instead of \\Fileserver1\Eng1, when Node1 fails and Node2 assumes ownership, the map will become invalid because the reference in the map is to \\Node1. If the map were created to the virtual name and Node1 were to fail, the map would still exist when the group associated with Eng1 failed over to Node2.

The previous diagram is an example and is not intended to imply limitations of a single group or node. Groups can contain multiple physical disks resources and file shares and nodes can have multiple groups, as shown by the group owned by Node2.

Cluster planning

Clustering servers greatly enhances the availability of file service by enabling file shares to fail over to additional storage servers, if problems arise. Clients see only a brief interruption of service as the file share resource transitions from one server node to the other.

Requirements for taking advantage of clustering include:

- Storage planning
- Network planning
- Protocol planning

Storage planning

For clustering, a storage unit (LUN) must be designated for the cluster and configured as a mirrorset. This LUN is used for the Quorum disk. The Quorum disk is the shared storage used by the cluster nodes to coordinate the internal cluster state.

One or more RAID arrays are dedicated to each cluster node for data storage. Each cluster node assumes ownership of at least one physical disk resource. That owner node serves all shares within that physical disks resource, until a failover condition occurs. When a failover occurs, the physical disk resource and all associated shares transition over to the remaining nodes and remain there until the other node is returned to service. Some types of shares are not cluster aware and are not available during a failover condition. See the “[Protocol planning](#)” section for additional information.

To prepare a basic disk for use in a cluster, a cluster group for each basic disk should be created to allow each resource to fail over separately. After the group is created, a physical disk resource is created in each of the groups. Cluster groups can contain more than one physical disk depending on the site-specific requirements. This physical disk resource is required for the basic disk to successfully work in a cluster environment, protecting it from simultaneous access from each node.



NOTE:

The LUN underlying the basic disk should be presented to only one node of the cluster using selective storage presentation SAN switch zoning, or having only one node online at all times until the physical resource for the basic disk is established.

In preparing for the cluster installation:

- All software components listed in the SAN Connection Guide must be installed and the fiber cables attached to the HBA(s) before the cluster installation is started.
- All shared disks, including the Quorum disk, must be accessible from both nodes. When testing connectivity between server and LUN, only one server should be given access to the LUN at a time or the non-testing server should be powered off.
- All shared disks must be configured as basic (not dynamic).
- All partitions on the disks must be formatted as NTFS.

Network planning

Clusters require more sophisticated networking arrangements than a stand alone storage server. For example, because a cluster must be deployed into a domain environment, workgroups are not supported. A Windows NT domain or Active Directory domain must be in place to contain the cluster names, virtual server names, and user and group information. A cluster cannot be deployed into a non domain environment.

All cluster deployments have at least seven network addresses and network names:

- The cluster name (Unique NETBIOS Name) and IP address
- Node A's name and IP address
- Node B's name and IP address
- At least one virtual server name and IP address for Node A
- At least one virtual server name and IP address for Node B
- Cluster Interconnect static IP addresses for Node A and Node B

In multi node deployments additional network addresses are required. For each additional node, three static IP addresses are required.

Virtual names and addresses are the only identification used by clients on the network. Because the names and addresses are virtual, their ownership can transition from one node to the other during a failover, preserving access to the shares on the virtual disks.

In addition, a cluster uses at least two network connections on each node:

- The cluster interconnect or “heartbeat” crossover cable connects to the first network port on each cluster node. In more than two node deployments, a private VLAN on a switch or hub is required for the cluster interconnect.
- The client network subnet connects to a second network port on each cluster node. The cluster node names and virtual server names have IP addresses residing on these subnets.



NOTE:

If the share is to remain available during a failover, each cluster node must be connected to the same network subnet. It is impossible for a cluster node to serve the data to a network to which it is not connected.

Protocol planning

The storage servers support many file sharing protocols, including sharing protocols for Windows, UNIX, Linux, Novell, Macintosh, Web, and FTP clients. However, not all of these protocols can take advantage of clustering. If a protocol does not support clustering, the share is not available to the clients until the owner cluster node is brought back online.

HP recommends placing cluster aware and non cluster aware protocols on different file shares.

Use the information in [Table 15](#) to determine whether it is advantageous to use clustering.

Table 15 Sharing protocol cluster support

Protocol	Client Variant	Cluster Aware (supports failover)	Supported
CIFS/SMB	Windows NT	Yes	Yes
	Windows 2000		
	Windows 95		
	Windows 98		
	Windows ME		
NFS	UNIX Linux	Yes	Yes
HTTP	Web	No	Yes
FTP	Many	Yes	Yes
NCP	Novell	No	Yes
AppleTalk	Apple	No	No



NOTE:

AppleTalk is not supported on clustered disk resources. AppleTalk requires local memory for volume indexing. On failover events, the memory map is lost and data corruption can occur.

Preparing for cluster installation

This section provides the steps necessary to cluster HP ProLiant storage servers.

Before beginning installation

Confirm that the following specifications have been met before proceeding:

- The SAN Connection Guide must be completed and all the necessary software components for connecting to the desired storage must be installed before the configuration of cluster services.
- It is required that at least one LUN has been presented for the configuration of the Quorum disk. This LUN must be created from shared storage and must be at least 50 MB. (500 MB is recommended). Additional LUNS may also be presented for use as shared disk resources.
- Cluster configurations should be deployed with dual data paths for high availability. Dual data paths from each node enable a path failure to occur that does not force the failover of the node. Clusters can be configured with single path, but if a failure in the path does occur, all of the node resources will be failed to the non-affected node.

Using Secure Path

Pathing software is required in configurations where multipathing to the storage is desired or required. For clustered products HP recommends maintaining two paths to the data as pathing software allows for datapath failure to occur without forcing a node failover. Secure Path is fully licensed and in contained

in select models. Secure Path is installed using the SAN Connection Guide, found on the desktop of the storage server.

Uninstalling Storage Manager



CAUTION:

The version of Storage Manager (Quota Management software located under the Shares tab) that comes installed on the server is not supported in a cluster. It must be uninstalled prior to creating the cluster. The uninstall tool is located in the Cluster Installation Guide under the Cluster tab in the WebUI. (See [Figure 93](#)).

The screenshot shows a Microsoft Internet Explorer window titled "Cluster Installation Guide - Microsoft Internet Explorer". The page content is as follows:

Cluster Installation Guide

Introduction

[Print this Cluster Install Guide](#)

The SAN connection tool must be completed and all the necessary software components for connecting to the desired storage must be installed before the configuration of cluster services.

It is required that at least one LUN has been presented for the configuration of the Quorum disk. This LUN must be created from shared storage and must be at least 50 MB (500MB is recommended). Additional LUNS may also be presented for use as shared disk resources.

Software Updates

Important: The version of Storage Manager (Quota Management software located under the *Shares* tab) that comes installed is not supported in a cluster. It must be uninstalled prior to creating the cluster.

[Click here to uninstall Storage Manager](#)

Once cluster installation has been completed, run the cluster update tool located at the [end of this document](#). This must be done after the Cluster install is complete. It is necessary to run this tool on **all cluster nodes**.

Checklists for Cluster Server Installation

This checklist assists in preparing for installation. Step-by-step instructions begin after the checklist.

Network Requirements

- A unique NetBIOS cluster name.
- Five unique, static IP addresses: two for the network adapters on the private network, two for the network adapters on the public network, and one for the cluster itself.
- A domain user account for Cluster service (all nodes must be members of the same domain).
- Each node should have at least two network adapters—one for connection to the public network and the other for the node-to-node private cluster network. If only one network adapter is used for both connections, the configuration is unsupported. A separate private network adapter is required for HCL certification.

Shared Disk Requirements:

Note: Do not let more than one node access the shared storage devices at the same time until Cluster service is installed on at least one node and that node is online. This can be accomplished through selective storage presentation, SAN switch zoning or having only one node online at all times.

Figure 93 Uninstall Storage Manager

Checklists for cluster server installation

These checklists assist in preparing for installation. Step-by-step instructions begin after the checklists.

Network requirements

- A unique NetBIOS cluster name.
- For each node deployed in the cluster the following static IP addresses are required:
 - One for the network adapters on the private network.
 - One for the network adapters on the public network.
 - One for the virtual server itself.

A single static cluster IP address is required for the entire cluster.

- A domain user account for Cluster service (all nodes must be members of the same domain).
- Each node should have at least two network adapters—one for connection to the public network and the other for the node-to-node private cluster network. If only one network adapter is used for both connections, the configuration is unsupported. A separate private network adapter is required for HCL certification.

Shared disk requirements



NOTE:

Do not let more than one node access the shared storage devices at the same time until Cluster service is installed on at least one node and that node is online. This can be accomplished through selective storage presentation, SAN switch zoning, or having only one node online at all times.

- All software components listed in the SAN Connection Guide must be installed and the fiber cables attached to the HBA(s) before the cluster installation is started.
- All shared disks, including the quorum disk, must be accessible from all nodes.
- All shared disks must be configured as basic (not dynamic).
- All partitions on the disks must be formatted as NTFS.

Cluster installation

During the installation process, nodes are shut down and rebooted. These steps guarantee that the data on disks that are attached to the shared storage bus is not lost or corrupted. This can happen when multiple nodes try to simultaneously write to the same disk that is not yet protected by the cluster software.

Use [Table 16](#) to determine which nodes and storage devices should be presented during each step.

Table 16 Power sequencing for cluster installation

Step	Node 1	Additional Nodes	Storage	Comments
Setting Up Networks	On	On	Not Presented	Verify that all storage devices on the shared bus are not presented; Power on all nodes.
Setting up Shared Disks	On	Off	Presented	Shut down all nodes. Present the shared storage, then power on the first node.
Verifying Disk Configuration	Off	On	Presented	Shut down first node, power on next node. Repeat this process for all cluster nodes.
Configuring the First Node	On	Off	Presented	Shut down all nodes; power on the first node.
Configuring additional Nodes	On	On	Presented	Power on the next node after the first node is successfully configured. Complete this process for all cluster nodes.
Post-installation	On	On	Presented	At this point all cluster nodes should be on.

To configure the Cluster service on the storage server, an account must have administrative permissions on each node. All nodes must be member servers within the same domain. It is not acceptable to have a mix of domain controllers and member servers in a cluster.

Setting up networks

Each cluster node requires at least two network adapters—one to connect to a public network, and one to connect to a private network consisting of cluster nodes only.

The private network adapter establishes node-to-node communication, cluster status signals, and cluster management. Each node's public network adapter connects the cluster to the public network where clients reside.

Verify that all network connections are correct, with private network adapters connected to other private network adapters only, and public network adapters connected to the public network.

Configuring the private network adapter

The following procedures are Best Practices provided by Microsoft and should be configured on the private network adapter.

- On the **General** tab of the private network adapter, ensure that only TCP/IP is selected.
- Ensure that the **Register this connection's address in DNS** is not selected in the DNS tab under advanced settings for the private network adapter.
- Ensure that the **Link Speed and Duplex** is set to 100Mps/Full Duplex under the **Advanced** tab for the Ethernet card used for the private network adapter.
- In all cases, set static IP addresses for the private network connector.

Configuring the public network adapter

While the public network adapter's IP address can be automatically obtained if a DHCP server is available, this is not recommended for cluster nodes. HP strongly recommends setting static IP addresses for all network adapters in the cluster, both private and public. If IP addresses are obtained via DHCP,

access to cluster nodes could become unavailable if the DHCP server goes down. If DHCP must be used for the public network adapter, use long lease periods to assure that the dynamically assigned lease address remains valid even if the DHCP service is temporarily lost. Keep in mind that Cluster service recognizes only one network interface per subnet.

Renaming the Local Area Network icons

HP recommends changing the names of the network connections for clarity. The naming helps identify a network and correctly assign its role. For example, "Cluster interconnect" for the private network and "Public connection" for the public network.

Verifying connectivity and name resolution

To verify name resolution, ping each node from a client using the node's machine name instead of its IP number.

Verifying domain membership

All nodes in the cluster must be members of the same domain and able to access a domain controller and a DNS Server.

Setting up a cluster user account

The Cluster service requires a domain user account under which the Cluster service can run. This user account must be created before installing Cluster service, because setup requires a user name and password. This user account should not belong to a user on the domain. This user account will need to be granted administrator privileges.

About the Quorum disk

When configuring the Quorum disk only one node should be powered on. All other potential cluster nodes must be powered off.

The quorum disk is used to store cluster configuration database checkpoints and log files that help manage the cluster. The Quorum disk must be a shared disk resource. HP makes the following Quorum disk recommendations:



NOTE:

Use the WebUI Disks tab to configure the Quorum disk resource.

- Dedicate a separate disk resource for a Quorum disk. Because the failure of the Quorum disk would cause the entire cluster to fail, HP strongly recommends that the disk resource be a RAID 1 configuration.
- Create a small partition [A minimum of 50 megabytes (MB) to be used as a quorum disk. HP recommends a Quorum disk be 500 MB.]

During the Cluster service installation, a drive letter must be provided for the quorum disk. HP recommends the drive letter Q for the quorum disk. It is also helpful to label the volume Quorum.

Configuring shared disks

Use the WebUI to configure additional shared disk resources. Verify that all shared disks are formatted as NTFS and are designated as Basic.

Additional shared disk resources are automatically added into the cluster as physical disk resources during the installation of cluster services. Each physical disk resource will reside in its own cluster group.

Verifying disk access and functionality

Write a file to each shared disk resource to verify functionality.

At this time, shut down the first node, power on the next node and repeat the Verifying Disk Access and Functionality step above for all cluster nodes. When it has been verified that all nodes can read and write from the disks, turn off the cluster nodes and power on the first, and then continue with this guide.

Configuring cluster service software

Clustering is installed by default. It is necessary to configure the cluster by launching Cluster Administrator. Follow the steps in the next section to configure the cluster. It is possible to add seven additional cluster nodes for an eight node cluster. Refer to the associated Storage Array documentation to determine the number of cluster nodes that are supported by the specific array in use under Windows Storage Server 2003.

Creating a cluster

From the WebUI, click the **Cluster** tab:




Figure 94 Cluster tab

1. Click **Advanced Cluster Management** to launch a Remote Desktop session.
2. Log into the Remote Desktop session.
3. Click **OK** when the error message regarding a cluster failure is displayed.
4. Select **File > New > Cluster**.

5. In the Welcome to the new server cluster window click **Next**.
6. In the New Server Cluster Wizard window, select the domain in which the cluster will be created and enter the name for the cluster. Click **Next**.
7. In the New Server Cluster Wizard window, enter the computer name to be the first node in the cluster, and then click **Next**.

The next step runs a pre-configuration analysis. This procedure analyzes and verifies the hardware and software configuration and identifies potential problems. A comprehensive and easy-to-read report is created, listing any potential configuration issues before the cluster is created.

1. Select the details tab to see a list of the items analyzed and any potential issues there may be with the cluster configuration.
2. If there are any issues fix the issues with the suggestion provided in the **Details** tab, and then click **Re-analyze**.

Some issues that can occur are:

- No shared disk for the Quorum disk. A shared disk must be created with a NTFS partition at least 50 MB in size.
- Use of DHCP addresses for network connections. All Network adapters must be configured with static IP addresses in a cluster configuration.
- File Services for Macintosh and Service for NetWare are not supported in a cluster configuration.
- Dynamic Disks are not supported in a cluster configuration.
- Errors appear on a network adapter that is not configured or does not have an active link. If the network adapter is not going to be used it should be disabled.

3. After all issues have been resolved, click **Next** to continue.
4. In the New Server Cluster Wizard window, enter the Cluster IP address, and then click **Next**.
5. In the New Server Cluster Wizard window, enter the login information for the domain account under which the cluster service will be run. Click **Next**.
6. In the New Server Cluster Wizard window, review the proposed cluster configuration and click **Next**.



NOTE:

It is possible to change the Quorum disk by clicking the Quorum button. This displays a list of available disks that can be used for the Quorum disk. Select the appropriate disk, and then click OK to continue.

7. In the New Server Cluster Wizard window, click **Next** to create the cluster.

After configuration is complete the following message is displayed: You have successfully completed the New Server Cluster Wizard.

8. Click **Finish** to close the wizard.

Adding nodes to a cluster



NOTE:

Only the Quorum disk should be accessible by the new node. The new node should not have access to the other LUNs in the cluster until after it has joined the cluster. After the node has joined the cluster, the LUNs may be presented to the new node. Move the physical disk resources over to the new node to confirm functionality.

1. Connect to the WebUI of a node that is a member of the cluster. Click the **Cluster** tab, and then click **Cluster Nodes**.
2. Ensure that the additional node has access to only the quorum LUN utilized as the cluster quorum disk.



CAUTION:

Presenting other LUNs to the non-clustered system could lead to data corruption.

3. Click **Add New Node**.
4. Enter the name of the node and specify the password for the cluster service account. Click **OK** to continue.




Figure 95 Adding a new node

5. In the Add Nodes Wizard window, click **Next**.
6. Specify the domain, and then click **Next**.
7. In the Add Nodes Wizard window, confirm the name of the node joining the cluster, and then click **Next**.
8. The next screen analyzes the configuration to determine the cluster configuration. Potential configuration errors are displayed. Fix any potential errors, and then click **Re-analyze**. Click **Next**.

9. In the Add Nodes Wizard window, enter the password for the cluster account, and then click **Next**.
10. The next screen displays a proposed cluster configuration summary. Confirm that all settings are correct, and then click **Next** to join the cluster.
11. Click **Next**, and then **Finish** to complete the cluster wizard.

After the node has successfully joined the cluster present all additional storage LUNS to the node. Please refer to the Cluster Storage Guide located in the **Cluster** tab for additional information on adding and configuring additional physical disk resources.

Geographically dispersed clusters

Cluster nodes can be geographically dispersed to provide an additional layer of fault tolerance. Geographically dispersed clusters are also referred to as stretched clusters.

The following rules must be followed with geographically dispersed clusters:

- A VLAN connection with latency of 500 milliseconds or less ensures that cluster consistency can be maintained. If the VLAN latency is over 500 milliseconds, the cluster consistency cannot be easily maintained.
- All nodes must be on the same subnet.

HP ProLiant Storage Server software updates

After cluster installation has been completed, run the cluster update tool located in the Cluster Installation Guide in the WebUI. The Cluster Installation Guide is located in the **Cluster** tab. This must be done after the cluster installation is complete. It is necessary to run this tool on all cluster nodes.




Figure 96 Cluster update tool

This completes the initial cluster installation.

Cluster groups and resources, including file shares

Management tasks for a cluster include creating and managing cluster resources and cluster groups. The Cluster Administrator tool provides complete online help for all cluster administration activities. Cluster resources are created and then assigned to logical, organizational groups. Ownership of these groups should be assigned in a balanced arrangement between the server nodes, distributing the processing load between the two nodes.

Cluster resources include administrative types of resources as well as file shares. The following paragraphs include overview and planning issues for cluster groups, cluster resources, and clustered file shares.

Creating and managing these resources and groups must be managed through Cluster Administrator, available from the **Cluster** tab of the WebUI. Complete online help for creating the various cluster objects is available in the Cluster Administrator tool.

Cluster group overview

A default cluster group is automatically created when the cluster is first created. This default cluster group contains an Internet Protocol (IP) Address resource, a Network Name resource, and the Quorum disk resource. When the new cluster is created, the (IP) address and the cluster name that were specified during setup are set up as the IP address and network name of this default cluster group.



CAUTION:

Do not delete or rename the Cluster Group or IP Address. Doing so results in losing the cluster and requires reinstallation of the cluster.

When creating groups, the administrator's first priority is to gain an understanding of how to manage the groups and their resources. Administrators may choose to create a resource group and a virtual server (IP Address resource and Network Name resource) for each node that will contain all resources owned by that node, or the administrator may choose to create a resource group and virtual server for each physical disk resource. Additionally, the administrator should try to balance the load of the groups and their resources on the cluster between the two nodes.

Node-based cluster groups

Creating only one resource group and one virtual server for each node facilitates group and resource administration. This setup allows administrators to include all file share resources under one group. Clients access all of the resources owned by one node through a virtual server name.

In node-based cluster groups, each group has its own network name and IP address. The administrator decides on which node to place each physical disk resource. This configuration provides a very coarse level of granularity. All resources within a group must remain on the same node. Only two IP addresses and network names are required. This configuration creates less overhead for resource and network administration. A possible disadvantage of this approach is that the resource groups can potentially grow large when many file shares are created.

Load balancing

The creation of separate cluster groups for each virtual server provides more flexibility in balancing the processing load on the cluster between the two nodes. Each cluster group can be assigned to a cluster node with the preferred owner parameter. For example, if there are two cluster groups, the cluster could be set up to have the first cluster group owned by node A and the second cluster group owned by node

B. This allows the network load to be handled by both devices simultaneously. If only one cluster group exists, it can only be owned by one node and the other node would not serve any network traffic.

Cluster resource overview

Hardware and software components that are managed by the cluster service are called cluster resources.

Resources represent individual system components. These resources are then organized into groups and managed as a group.

Some resources are created automatically by the system and other resources must be set up manually.

Resource types:

- IP Address resource
- Cluster name resource
- Cluster Quorum disk resource
- Physical Disk resource
- Virtual server name resources
- CIFS file share resources
- NFS file share resources

File share resource planning issues

CIFS and NFS are cluster aware protocols that support the Active/Active cluster model, allowing resources to be distributed and processed on both nodes at the same time. For example, some NFS file share resources can be assigned to a group owned by a virtual server for NodeA and additional NFS file share resources can be assigned to a group owned by a virtual server for NodeB.

Configuring the file shares as cluster resources provides for high availability of file shares. Because the resources are placed into groups, ownership of the files can easily move from one node to the other, as circumstances require. If the cluster node owning the group of file shares should be shut down or fail, the other node in the cluster will begin sharing the directories until the original owner node is brought back on line. At that time, ownership of the group and its resources can be brought back to the original owner node.

Resource planning

1. Create at least one virtual server for each node in the cluster.

A virtual server is a resource group consisting of an IP Address resource and a Network Name resource. Ownership of these virtual servers should be assigned to the different server nodes. In addition to providing load balancing capabilities, the virtual server allows for the transition of group resources in failover situations.

2. Create a virtual server group for each node in the cluster.

Cluster resource groups are used to balance the processing load on the servers. Distribute ownership of the groups between the virtual servers.

3. For NFS environments, configure the NFS server.

NFS specific procedures include entering audit and file lock information as well as setting up client groups and user name mappings. These procedures are not unique to a clustered deployment and are detailed in the “Microsoft Services for NFS” chapter. Changes to NFS setup information are automatically replicated to all nodes in a cluster.

4. Create the file share resources.

In a clustered environment, file shares are created as a type of cluster resource. Creating cluster resources and file shares is documented later in this chapter.

5. Assign ownership of the file share resources to the resource groups.

- a.** Divide ownership of the file share resource between the resource groups, which are in turn distributed between the virtual servers, for effective load balancing.
- b.** Verify that the physical disk resource for this file share is also included in this group.
- c.** Verify that the resources are dependent on the virtual servers and physical disk resources from which the file share was created.

Permissions and access rights on share resources

File Share and NFS Share permissions must be managed via the Cluster Administrator tool versus the individual shares on the file system themselves via Windows Explorer. Administering them through the Cluster Administrator tool allows the permissions to migrate from one node to other. In addition, permissions established using Explorer are lost after the share is failed or taken offline. To access the permissions, see “[Setting permissions for an SMB file share](#)” and “[Setting permissions for an NFS share](#).”

NFS cluster-specific issues

In addition to the user name mapping best practices outlined in the “[Services for NFS/UNIX](#)” chapter, there are additional recommendations.

For convenience, all suggestions are listed below:

- Back up user and group mappings
To avoid loss of complex advanced mappings in the case of a system failure, back up the mappings whenever the mappings have been edited or new mappings have been added.
- Map consistently
Groups that are mapped to each other should contain the same users and the members of the groups should be properly mapped to each other to ensure proper file access.
- Map properly
 - Valid UNIX users should be mapped to valid Windows users.
 - Valid UNIX groups should be mapped to valid Windows groups.
 - Mapped Windows user must have the **Access this computer from the Network privilege** or the mapping will be squashed.
 - The mapped Windows user must have an active password, or the mapping will be squashed.
- In a clustered deployment, create user name mappings using domain user accounts.
Because the security identifiers of local accounts are recognized only by the local server, other nodes in the cluster will not be able to resolve those accounts during a failover. Do not create mappings using local user and group accounts.
- In a clustered deployment, administer user name mapping on a computer that belongs to a trusted domain.
If NFS administration tasks are performed on a computer that belongs to a domain that is not trusted by the domain of the cluster, the changes are not properly replicated among the nodes in the cluster.
- In a clustered deployment, if PCNFS password and group files are being used to provide user and group information, these files must be located on each node of the system.
Example: If the password and group files are located at *c:\maps* on node 1, then they must also be at *c:\maps* on node 2. The contents of the password and group files must be the same on both nodes as well.

These password and group files on each server node must be updated periodically to maintain consistency and prevent users or groups from being inadvertently squashed.

Non cluster aware file sharing protocols

Services for Macintosh (SFM), File and Print Services for NetWare, HTTP file sharing protocols are not cluster aware and will experience service interruption if installed on a clustered resource during failover events of the resource. Service interruptions will be similar to those experienced during a server outage. Data that has not been saved to disk prior to the outage will experience data loss. In the case of SFM, it is not supported because SFM maintains state information in memory. Specifically, the Macintosh volume index is located in paged pool memory. Using SFM in clustered mode is not supported and may result in data loss similar in nature to a downed server should the resource it is based on fails over to the opposing node.

Creating a new cluster group

To create a cluster group:

1. Open the WebUI to the **Cluster** tab, and then click **Cluster Groups**.

Group	State	Owner	Description	Group Tasks
Cluster Group	Online	DOCBOX9000		[Move] [Take Offline] [Bring Online] [Delete] [Create] [Properties] [Advanced Cluster Management]
Group 0	Online	DOCBOX9000		
Group 3	Online	DOCBOX9000		
Group 4	Online	DOCBOX9000		
group 44	Online	DOCBOX9000		
Group 54	Online	DOCBOX9000		
Printer Group	Online	DOCBOX9000		

Figure 97 Cluster Groups page

2. Click **Create** to create a new group.
3. Specify the properties for the new cluster group, and then click **OK** to create the cluster group.

Adding new storage to a cluster

Present the new storage to one node in the cluster. This can be accomplished through selective storage presentation or through SAN switch zoning.

1. Open the WebUI and navigate to the **Disks** tab.
2. Click the **Disks** subtab.
3. Select the disk that needs to be configured from the list of available disks, and then click **Create New Volume**.
4. Follow the steps in the wizard to create the new volume. The LUN needs to be configured as a basic disk with a NTFS file system.



NOTE:

If the disk does not appear in the list of available disks on the Manage Disks page then click Rescan to rescan for new disks and refresh the page.

5. Open the WebUI and click the **Cluster** tab.
6. Follow the procedures listed below to create a physical disk resource.

See the Cluster Storage Guide on the WebUI **Cluster** tab for detailed information on adding storage elements into the cluster.

Creating physical disk resources

A physical disk resource must reside within a cluster group. An existing cluster group can be used or a new cluster group must be created. See “[Creating a new cluster group](#)” earlier in this chapter.

To create a physical disk resource:

1. On the **Cluster** tab, click **Cluster Resources**.




Figure 98 Cluster Resources page

2. Click **Create**.
3. On the Welcome Page click **Next**.
4. Specify a name for the cluster resource and enter a description for the resource.
5. Select the Cluster group the physical disk will reside in.
6. Select Physical Disk as the resource type, and then click **Next**.
7. Select the Possible Owners, and then click **Next**.
8. Set the dependencies, and then click **Next**.



NOTE:

Physical disk resources usually do not have any dependencies set.

9. Specify the available disk resource, and then click **Next**.
10. Review the configuration, and then click **Finish** to create the physical disk resource.
11. After the resource is created it is necessary to bring it online. In the **Cluster Resources** page, select the resource and click **Bring Online**.
12. Click **OK** on the Bring a Resource Online page to bring the new physical disk resource online.

13. Present the LUN to the additional cluster nodes.
14. Move the physical disk resource to the other nodes to confirm functionality.

To move a resource:

1. On the **Cluster** tab in the WebUI, click **Cluster Groups**.
2. Select the group, and then click **Move**.
3. Specify the new location for the group, and then click **OK**.



NOTE:

In multi-node clusters it is necessary to specify the node to move the group to. When a cluster group is moved to another node all resources in that group are moved.



NOTE:

When a physical disk resource is owned by a node, the disk appears as an unknown unreadable disk to all other cluster nodes. This is a normal condition. When the physical disk resource moves to another node, the disk resource then becomes readable.

Creating file share resources

To create a file share resource:

1. Open the WebUI to the **Cluster** tab, and then click **Cluster Resources**.
2. Click **Create**.
3. Click **Next** to proceed with the wizard.
4. Specify the share name, description, the cluster group in which the share will reside, and the resource type. The resource type should be File Share.



NOTE:

A file share resource must reside in the same cluster group as the physical disk resource it will reside on.




Figure 99 Creating a file share resource

5. Click **Next**.
6. Select the possible owners, and then click **Next** to continue.
7. Specify the dependencies for the file share resource. The physical disk should always be selected as a dependency. It is also possible to set the file share to be dependent on the network name for the cluster group in which it will reside. If a dependency is set on the network name the file share resource will not come online until the network name resource does. Click **Next** to continue.



NOTE:

The physical disk resource specified in this step must reside in the same cluster group as specified in the beginning of this wizard.

8. Specify the share name, path, and user limit, and then click **Next** to continue.
9. Review the configuration, and then click **Finish** to create the share.
10. After the resource is created it is necessary to bring it online. In the Cluster Resources page, select the resource, and then click **Bring Online**.
11. Click **OK** on the **Bring a Resource Online** page to bring the new file share resource online.

Setting permissions for an SMB file share

When a share resource is created via the WebUI and brought online, the default permission is set to: Everyone=Read-Only.

To change the default permissions:

1. On the **Cluster** tab, click **Advanced Cluster Management**.
2. Log into Remote Desktop.
3. Click the group.

4. Right-click the resource, and then click **Properties**.




Figure 100 Resource parameters for SMB file share

5. Click the **Parameters** tab.
6. Click **Permissions**.




Figure 101 Set resource permissions

7. Set the permissions, and then click **OK**.

Creating NFS share resources

To create an NFS share resource:

1. Open the WebUI to the **Cluster** tab, and then click **Cluster Resources**.
2. Click **Create**.
3. Click **Next** to proceed with the wizard.
4. Specify the name, description, the cluster group in which the share will reside, and the resource type. The resource type should be NFS Share.
5. Click **Next**.
6. Select the possible owners, and then click **Next** to continue.
7. Specify the dependencies, and then click **Next** to continue.
8. Specify the share name, path, Share Root Only or Share Sub-directories only, encoding, anonymous access, and anonymous UID/GID, and then click **Next** to continue.
9. Review the configuration, and then click **Finish** to create the NFS share.
10. After the resource is created it is necessary to bring it online. On the **Cluster Resources** page, select the resource, and then click **Bring Online**.

11. Click **OK** on the Bring a Resource Online page to bring the new resource online.

Setting permissions for an NFS share

When a share resource is created via the WebUI and brought online, the default permission is set to: Everyone=Read-Only.

To change the default permissions:

1. On the **Cluster** tab, click **Advanced Cluster Management**.
2. Log into Remote Desktop.
3. Click the group.
4. Right-click the resource, and then click **Properties**.




Figure 102 NFS Share Resource parameters

5. Click the **Parameters** tab.
6. Click **Permissions**.




Figure 103 Set NFS Share resource permissions

7. Set the permissions, and then click **OK**.

Creating IP address resources

1. Open the WebUI to the **Cluster** tab, and then click **Cluster Resources**.
2. Click **Create**.
3. Click **Next**.
4. Specify the name, description, the cluster group in which the resource will reside, and the resource type. The resource type should be IP Address.
5. Click **Next**.
6. Click the possible owners, and then click **Next** to continue.
7. Specify the dependencies, and then click **Next** to continue.
8. Specify the address, subnet mask, network and whether to enable NetBIOS for this address, and then click **Next**.




Figure 104 Creating an IP address resource

9. Review the configuration and click **Finish** to create the resource.
10. After the resource is created it is necessary to bring it online. In the Cluster Resources page, select the resource, and then click **Bring Online**.
11. Click **OK** on the Bring a Resource Online page to bring the new resource online.

Creating network name resources

1. Open the WebUI to the **Cluster** tab, and then click **Cluster Resources**.
2. Click **Create**.
3. Click **Next** to proceed with the wizard.
4. Specify the name, description, the cluster group in which the resource will reside, and the resource type. The resource type should be Network Name.
5. Click **Next**.
6. Select the possible owners, and then click **Next**.
7. Specify the dependencies, and then click **Next**.



NOTE:

The resource must be dependent on an IP address resource.




Figure 105 Network Name Parameters

8. Select whether or not DNS registration must succeed and whether to enable kerberos authentication and click **Next** to continue.
9. Review the configuration, and then click **Finish** to create the resource.
10. After the resource is created it is necessary to bring it online. On the Cluster Resources page, select the resource and click **Bring Online**.
11. Click **OK** on the **Bring a Resource Online** page to bring the new resource online.

Basic cluster administration procedures

Failing over and failing back

As previously mentioned, when a node goes offline, all resources dependent on that node are automatically failed over to another node. Processing continues, but in a reduced manner because all operations must be processed on the remaining node(s). In clusters containing more than two nodes, additional fail over rules can be applied. For instance, groups can be configured to fail over different nodes to balance the additional work load imposed by the failed node. Nodes can be excluded from the possible owners list to prevent a resource from coming online on a particular node. Lastly the preferred owners list can be ordered, to provide an ordered list of failover nodes. Using these tools, the failover of resources can be controlled with in a multinode cluster to provide a controlled balanced failover methodology that balances the increased work load.

Because operating environments differ, the administrator must indicate whether the system will automatically fail the resources (organized by resource groups) back to their original node or will leave the resources failed over, waiting for the resources to be moved back manually.



NOTE:

If the storage server is not set to automatically fail back the resources to their designated owner, the resources must be moved back manually each time a failover occurs.

Restarting one cluster node



CAUTION:

Restarting a cluster node should be done only after confirming that the other node(s) in the cluster are functioning normally. Adequate warning should be given to users connected to resources of the node being restarted. Attached connections can be viewed through the Management Console on the storage server Desktop using Terminal Services. From the Management Console, select **File Sharing > Shared Folders > Sessions**.

The physical process of restarting one of the nodes of a cluster is the same as restarting a storage server in single node environment. However, additional caution is needed.

Restarting a cluster node causes all file shares served by that node to fail over to the another node(s) in the cluster based on the failover policy in place. Until the failover process completes, any currently executing read and write operations will fail. Other node(s) in the cluster will be placed under a heavier load by the extra work until the restarted node comes up and the resources are moved back.

Shutting down one cluster node



CAUTION:

Shutting down a cluster node must be done only after confirming that the other node(s) in the cluster are functioning normally. Adequate warning should be given to users connected to resources of the node being shutdown.

Shutting down a cluster node causes file shares served by that node to fail over to the other node(s). This causes any currently executing client read and write operations to fail until the cluster failover process completes. The other node(s) are placed under a heavier load by the extra work until the second node is powered up and rejoins the cluster.

Powering down the cluster

The power down process for the storage server cluster is similar to the process for a single node, but with the cluster, extra care must be taken with the storage subsystem and the sequence of the shutdown.

The power down process is divided into two main steps:

1. Shutting down the cluster nodes
2. Removing power from the cluster nodes

The sequence of these steps is critical. The devices must be shut down before the storage subsystem. Improperly shutting down the nodes and the storage subsystem causes corruption and loss of data.



CAUTION:

Before powering down the cluster nodes, follow the proper shutdown procedure as previously illustrated. See "["Shutting down one cluster node."](#)" Only one cluster node should be shut down at a time.



CAUTION:

The cluster nodes should never be powered on when the storage subsystem is not available.

Powering up the cluster

The power up process for the storage server cluster is more complex than it is for a single node because extra care must be taken with the storage subsystem.

The sequence of the power up steps is critical. Improper power up procedures can cause corruption and loss of data.



CAUTION:

Do not power up the cluster nodes without first powering up the storage subsystem, and verifying it is operating normally.

Nodes should be powered up separately allowing one node to form the cluster prior to powering up the additional node(s). To power up the cluster nodes:

1. After the storage subsystem is confirmed to be operating normally, power up a single node. Wait for the node to come completely up before powering up the subsequent node.

If more than one node is powered up at the same time, the first node that completes the sequence gains ownership of the cluster quorum and controls the cluster database. Designate a particular node as the usual cluster quorum owner by always powering up that node first and letting it completely restart before powering up additional cluster node(s).

2. Power up the additional cluster node(s). Each node should be allowed to start fully, prior to starting a subsequent node.

As each node starts, the monitor displays the logon dialog. Background processes start the cluster service and form the cluster.

Shadow copies in a clustered environment

The creation and management of clustered Shadow Copy resources in a cluster should be performed using the WebUI by clicking the **Disk** tab, and then clicking **Shadow Copy** or from the file system by right-clicking the volume, and then clicking **Shadow Copy**.

Assuming the underlying disk is part of a cluster, both methods will generate a cluster resource on the cluster that is viewable from Cluster Administrator and the **Cluster Resource** tab of the WebUI. While the ability to create the Shadow Copy Resource is available in the Cluster Administrator Management Tool, this operation is not supported by Microsoft. The resource may be taken offline/online and managed with the group via all means available.

As recommended in the Shadow Copy chapter, the location of the cache file is recommended on a separate disk from the original data. In this case, a physical disk resource for the cache file disk should be created in the same cluster group as the intended Shadow Copy resource and the volume for which snapshots will be enabled. The resource should be created prior to the establishment of Shadow Copies. The Shadow Copy resource should be dependent on both the original physical disk resource and the physical disk resource that contains the cache file. The update of the Shadow Copy schedule may be done via the Cluster Administrator tool, the WebUI, or the file system.

Creating a cluster printer spooler

Printer spoolers should be created in a separate group dedicated to this purpose for ease of management. For each printer spooler a physical resource is required to instantiate the print spooler resource. In some cases, dedicated physical resources are not available and hence sharing of the physical resource among other members of the group is acceptable, remembering that all members of a group are managed as a unit. Hence, the group will failover and fallback as a group.

To create a printer spooler:

1. Create a dedicated group (if desired).
2. Create a physical resource (disk) (if required, see note).
3. Create an IP address resource for the Virtual Server to be created (if required, see note).
4. Create a Virtual Server Resource (Network Name) (if required, see note).



NOTE:

Steps 1-4 may be done via the WebUI interface using the appropriate functions or via the Advanced Cluster Management function and are documented elsewhere in this chapter. If the printer spool resource is added to an existing group with a physical resource, IP address and virtual server resource, steps 1-4 are not required.

5. Create a Print Spool resource:
 - a. Click the **Cluster** tab.
 - b. Click **Advanced Cluster Management**.
 - c. Select the group container for the printer spooler.
 - d. Right-click and click **Printer Resource**.
 - e. Enter the name of the printer resource.
 - f. Select all of the appropriate dependent resources (IP address, network name, and physical resource).
 - g. Select the folder to place the spooler temporary contents, and then click **Finish**.
 - h. Close Cluster Administrator.
6. To connect to the Virtual Server Name or IP address created in the steps above:
Select **Start > Run**, and then type:
`\\"virtual_server_name or ip address"` from the local menu
A session opens to the virtual server.
7. To add a printer to the virtual server:
 - a. Double-click the printers and faxes icon.
 - b. Right-click the new screen, and then click **add printer**. A wizard starts.
 - c. Click **create a new port**, and then click **Next**.
 - d. Enter the IP address of the network printer.
 - e. Update the Port Name if desired, click **Next**, and then click **Finish**.

- f. Select the appropriate driver, and then click **Next**.
- g. If presented with a dialog to replace the driver present, click **keep the driver**, and then click **Next**.
- h. Name the printer, and then click **Next**.
- i. Provide a share name for the printer for network access, and then click **Next**.
- j. Provide location information and comments, and then click **Next**.
- k. Click **Yes** to print a test page, click **Next**, and then click **Finish**.
- l. A dialog box appears regarding the test page. Select the appropriate answer.

The Printer Spool is now a clustered resource.

A NIC Teaming

Some models of the HP ProLiant Storage Server are equipped with the HP Network Teaming and Configuration utility. The utility allows administrators to configure and monitor Ethernet network interface controllers (NIC) teams in a Windows-based operating system. These teams provide options for increasing fault tolerance and throughput.

Fault tolerance provides automatic redundancy. If the primary NIC fails, the secondary NIC takes over. Load Balancing provides the ability to balance transmissions across NICs.



NOTE:

Select models ship with the NIC teaming utility available, however it must be installed and configured.



NOTE:

Installing NIC teaming requires a restart of the server.

Installing the HP Network Teaming Utility

Before using the HP Network Teaming utility, it must be installed.



NOTE:

Installing and configuring NIC teaming should always be performed via the iLO port or the console using a direct attached keyboard, monitor, and mouse since IP connections could be reset during the configuration process. Do not use Remote Desktop.

To install the HP Network Teaming utility:

1. In the URL field of the Web browser, enter the IP address of the Integrated Lights-Out port.



NOTE:

The iLO port requires a license key. The key is included with the product inside the Country Kit. Refer to the iLO Advanced License Pack for activation instructions.

**NOTE:**

The iLO port can also be accessed from the HP Utilities tab of the WebUI by clicking the remote management link.

2. At the Integrated Lights-Out Account Login window, supply the username and password for the iLO, and then click **Login**.
3. Click the **Remote Console** tab.
4. Click on the **Remote Console** choice in the menu on the left side of the screen.
5. Press **Ctrl-Alt-Del** to log into the console.
6. Supply an administrator username and password.
7. Double-click the **NIC Team Setup** icon on the desktop.
8. When the following message box is displayed, click **Install**.




Figure 106 Installing Network Teaming

9. When the installation process is complete, the following dialog box is displayed.




Figure 107 Network Teaming installation complete

10. Click **Close**.
11. Restart the system.

**CAUTION:**

To ensure proper functioning of the software, the server must be restarted at this time.

Opening the HP Network Teaming Utility

The HP Network Teaming utility is now accessible from the Windows toolbar at the bottom of the storage server desktop. To open the utility, click the **HP Network Teaming utility** icon.



Figure 108 HP Network Teaming utility icon

Adding and configuring NICs in a team

Before a NIC is teamed, verify the following:

- The NICs must be on the same network.
- The NICs must be DHCP enabled and the DNS server address must be left blank.

**NOTE:**

The teaming utility becomes unstable if static IP addresses, subnets, and DNS addresses are set before teaming.

- Duplex and speed settings must be set to use the default values.

To team the NICs:

1. Open the HP Network Teaming utility.

The **Network Teaming and Configuration Properties** dialog box is displayed. The type of NIC and the slot and port used is shown.




Figure 109 HP Network Teaming Properties dialog box

2. Highlight the NICs to team.
3. Click **Team**.




Figure 110 NIC Properties, Teaming Controls tab, Fault Tolerant option

4. Configure the team by choosing either **Fault Tolerant** or **Load Balancing**.

The fault tolerance and load balancing options are discussed in the following sections.

5. Click **OK** to accept the team properties.
6. Click **OK** in the HP Network Teaming and Configuration Properties dialog box to apply the changes.
7. Click **Yes** when prompted to apply all configuration changes. Wait while the adapters are configured. This process could take several seconds.
8. The following screen is displayed, indicating that there are additional procedures to perform in the NIC teaming process.




Figure 111 HP Network Teaming dialog box

9. Click **Yes** to reboot now.

Fault tolerance

The Fault Tolerance teaming option provides three redundancy control options:

- **Manual**—This setting allows change from a Primary NIC to a Secondary NIC only when **Switch Now** is clicked.



NOTE:

The **Switch Now** option is disabled until **Manual** is selected, and then **OK** is clicked.

- **Fail on Fault**—Automatically switches from a primary NIC to a secondary NIC when the primary NIC fails.
- **Smart Switch**—Allows a member of a team be selected as the preferred Primary Smart Switch NIC. As long as this NIC is operational, it is always the active NIC. If the NIC fails and it is eventually restored or replaced, it automatically resumes its status as the active NIC.



NOTE:

HP recommends **Smart Switch** for fault tolerance.

Detailed information about configuring teams for fault tolerance can be found in the HP Network Teaming Utility help.

Load balancing

The **Load Balancing** teaming option provides four load balancing control options:




Figure 112 NIC Properties, Teaming Controls tab, Load Balancing option

Detailed information about these four load balancing teaming options can be found in the HP Network Teaming help.

- **Transmit Load Balancing**—All transmit IP frames are load balanced among all team members based on a Load Balancing algorithm in the teaming device driver. The Current Primary adapter transmits all other frames, and receives all frames for the team. If a failover event occurs, one of the non-Primary adapters assumes the role of Current Primary adapter, and transmit IP packets are load balanced among all remaining team members. If a failure occurs in any of the non-Primary adapters, the packets are load balanced among all remaining team members.
- **Switch-assisted Load Balancing**—All transmit packets are load balanced among all team members based on a Load Balancing algorithm in the teaming device driver. The receive packets are load balanced among all team members by the switch. If a failure of any team member occurs, the packets are load balanced among the remaining adapters. There is no primary adapter in a Switch-assisted Load Balancing team.
- **Balance with MAC Address**—Allows load balancing of IP packets among the teamed NICs using the last four bits of the MAC Address. (See following Note.)
- **Balance with IP Address**—Allows load balancing of IP packets among the teamed NICs using the last four bits of the IP Address. (See following Note.)



NOTE:

The teaming utility can load balance IP packets among the teamed NICs installed in a server. The primary NIC in the team receives all incoming packets. The choice is available to load balance with the source MAC address (the address transmitted from the workstation) or the source IP address.

Using the last four bits of either source address, the teaming driver algorithm assigns this source address to the port of one of the NICs in the team. This port is then used to transmit all packets destined for that source address. If there are four NICs in the team, the packets are received by the primary NIC on the team. The packets are retransmitted through one of the four ports.

Configuring the NIC team properties

At this point, the NICs are teamed but are not completely configured. Additional procedures include:

- Renaming the teamed connection
- Selecting the option to show an icon on the taskbar
- Configuring TCP/IP on the new team

Renaming the teamed connection

The assigned name for the new NIC team connection is “Local Area Connection X,” where X represents the next available connection number generated by the system. HP recommends changing this name to a more meaningful name, such as “NIC Team.”

To change the name of the connection:

1. From the desktop, right-click the **My Network Places** icon, and then click **Properties**.
2. Move the cursor over each connection icon to view the pop up box of the icon's name. Locate **HP Network Teaming Virtual Miniport**.
3. Right-click the connection icon for **HP Network Teaming Virtual Miniport**, and then click **Rename**. Enter a name that is more descriptive than “Local Area Connection X,” such as “NIC Team.”

Showing a connection icon on the taskbar

To show a connection icon:

1. In the **Network and Dial up Connections** screen, double-click the **NIC Team** connection, and then click **Properties**.
2. At the bottom of the screen, select **Show icon in task bar when connected**, and then click **Close**.

Configuring the TCP/IP protocol on the new team

After teaming the NICs, a new virtual network adapter for the team is automatically created. However, by default the new adapter is set to DHCP. To manually configure the IP address, perform the following steps.

To enter the TCP/IP address information for the team:

1. From the desktop, go to the **Network and Dial up Connections** window, and then click **Properties**. Right-click the **NIC Team** icon, and then select **Properties**.




Figure 113 NIC Team Properties dialog box

2. Use the arrows and the scroll bar on the right of the screen to scroll through the **Components** list.
3. Click **Internet Protocol (TCP/IP)**, and then click **Properties**.




Figure 114 NIC Team TCP/IP Properties dialog box



NOTE:

If a NIC is teamed, do not modify the TCP/IP settings for the individual NIC ports.

4. Select **Use the following IP address**, and then enter the IP address and subnet mask. If desired, enter the default gateway.
5. Click **OK**. The Ethernet Team should be working.

Checking the status of the team

To check the status of the Ethernet Team, open the HP Network Teaming utility. The **Configuration Properties** window is displayed, showing the teamed NICs.




Figure 115 NIC Teaming status

NIC teaming troubleshooting

Problems with the NIC teaming feature are diagnosed by the connection icons displayed in the **HP Network Teaming and Configuration** dialog box. The following table lists the error icons for RJ 45 NICs.

Table 17 NIC Teaming Troubleshooting

RJ-45	Description
	Active OK—The NIC is operating properly. The driver is installed in the registry and is loaded. If the NIC is a member of a team, the NIC is active.
	Installed inactive—The NIC is installed and is OK, but is not active.
	Cable fault—The driver is installed in the registry and is loaded. The broken cable indicator means that the cable is unplugged, loose, broken, or the switch or hub is not operating properly. If this icon is displayed, check all network connections and make sure the hub/switch is working properly. When the connection is restored, this icon will change.
	Inactive cable fault—A cable fault has occurred while the NIC was inactive.
	Hardware failure—The driver is installed in the registry and is loaded. The driver is reporting a hardware problem with the NIC. This indicates a serious problem. Contact your HP authorized service provider.

RJ-45	Description
	Unknown—The server is unable to communicate with the driver for the installed NIC. The NIC is installed in the registry, but the driver is not. This error occurs when the NIC has been installed but the server has not been restarted. If this problem persists after the server has been restarted, the driver has not been loaded or the Advanced Network Control utility is unable to communicate with the driver. <i>Note: Only NICs assigned as members of a team are displayed as Unknown. If a teamed NIC is turned off, it displays as Unknown.</i>
	Disabled—The NIC has been disabled through the Device Manager or NCPA.

For more advanced problems with NIC teaming, refer to the help section in the HP Teaming and Configuration utility.

Index

A

access rights, managing, 189
ACL
 defined, 99
 translating, 146
AFP
 installing services for AppleTalk, 105
 installing services for Macintosh, 105
 shares, setting up, 106
alerts, e-mail, setting up, 28
audit logs, 25
Authentication software, installing, 130
authorized reseller, HP, 15

B

backup
 mappings, 152
 with shadow copies, 77

C

cache file, shadow copies, 67
CIFS
 share support, 100
CIFS/SMB
 administration, 79
client groups
 adding NFS, 142
 deleting NFS, 143
 editing NFS, 143
 managing NFS, 141
cluster
 adding new storage, 191
 analysis, 184
 components, hierarchy, 174
 concepts, 173
 concepts, diagram, 174
 creating, 183
 diagram, 172
 dual data paths, 178
 geographically dispersed, 186
 group, 187
 group, creating, 190
 groups, node-based, 187
 installation, 180
 installation checklist, 179
 load balancing, 188
 managing access rights, 189
 managing file share permissions, 189
 multi node support, 171
 network requirements, 180
 NFS issues, 189

nodes
 adding, 185
 powering down, 201
 powering up, 202
 restarting, 201
overview, 171, 171
planning, 175
preparing for installation, 178
printer spooler, 203
protocols, non cluster aware, 190
resource overview, 188
resources, 187
resources, defined, 172
setting up user account, 182
shared disk requirements, 180
configuring
 private network adapter, 181
 shared disks, 183
connectivity, verifying, 182
conventions
 document, 13
 text symbols, 13
creating NFS file shares, 133

D

date, system, changing, 24
deployment scenarios, 18
directory quotas, establishing, 116
disk access, verifying, 183
document
 conventions, 13
 prerequisites, 13
 related documentation, 13
domain environment, 20
domain membership, verifying, 182
dual data paths, 178

E

e-mail alerts, setting up, 28
encoding types, 137
environments
 overview, 19
Ethernet NIC teams
 adding, 207
 checking status, 214
 configuring, 209
 configuring properties, 212
 configuring TCP/IP, 212
 renaming the connection, 212
 setting up, 30, 205
 showing connection icon, 212

troubleshooting, 215
events, Services for NFS, logging, 128
Exchange Server, 19
explicit group mapping, 150
explicit mappings, 145, 149
exports, 126

F

fail on fault setting, 210
failover
 automatic, 200
 defined, 173
 resources, 173
fault tolerance
 for NIC teams, 210
File and Print Services for NetWare. See FPNW, 157
file level permissions, 92
file recovery, 75
file screening, 117
file server consolidation, 19
file share
 resource planning, 188
file share permissions, 194
file share permissions, managing, 189
file share resources, 175, 193
files, ownership, 97
folder recovery, 75
folders
 auditing access, 95
 compress tab, 89
 creating new, 88
 creating new share, 90
 deleting, 89
 managing, 87
 managing shares for, 91
 modifying properties, 89
 navigating to, 87
FPNW
 accessing, 158
 described, 157
 installing, 157

G

getting help, 14
group names
 examples, 80
 managing, 80
group, cluster
 cluster
 group, 175
groups
 adding from a domain, 86
 adding local users, 85
 adding to permissions list, 93
 local, adding, 83
 local, deleting, 84
 local, modifying properties, 84
 properties, general tab, 85

properties, members tab, 85
removing local users, 86

H

hardware features, 17
help, obtaining, 14
HP
 authorized reseller, 15
 storage web site, 15
 technical support, 14
 Web Jetadmin, 123
HP Network Teaming Utility
 installing, 205
 opening, 207

I

iLO. See Integrated Lights-Out Port, 166
Insight Manager
 described, 169
installation, cluster, preparing for, 178
Integrated Lights-Out port
 accessing storage servers, 168
 activating, 30
 configuration, 168
 described, 166
 features, 167
 license key, 30
IP address resource, 175, 198

J

Jetadmin, 123

L

LAN icons, renaming, 182
license key, iLO port, 30
load balancing, 188, 211
 switch-assisted, 211
 transmit, 211
 with IP address, 211
 with MAC address, 211
localhost, 127
locks, NFS, 140
logging, Services for NFS events, 128
logs
 accessing, 25
 audit, 25
 options, 25
LUNs
 presenting to cluster node, 186

M

Macintosh, installing services for, 105
Management Console, 22
managing system storage, 29
mappings

backup and restore, 152
best practices, 146
creating, 147
data stored, 148
explicit, 145, 149
NFS, 144
simple, 145, 148
squashed, 145
Microsoft Feature Pack, 19
mounted drives and shadow copies, 64
multiprotocol environments, 19

N

NCP

creating new share, 162, 164

NetWare

adding local users, 160
enabling user accounts, 160
installing services for, 157
supervisor account, 161
network name resource, 175, 199
network planning, 176
network requirements, cluster, 180
network settings, changing, 28
networks
 setting up, 181

NFS

async/sync settings, 139
authenticating user access, 126
client groups, 141
 adding, 142
 deleting, 143
 editing, 143
cluster specific issues, 189
compatibility issues, 100
deleting shares, 135
file share, creating, 133
file shares, creating, 133
file sharing tests, 153
group mappings, 144
locks, 140
modifying share properties, 135
protocol properties settings, 138
Server settings, 129
share properties, 139
user mapping server, 127
 user mappings, 144
NFS only access, 138
NFS share permissions, 197
NFS share resource, 196
node
 defined, 172
NTFS permissions, 132

P

passwords
 modifying local userxd5 s, 82
permissions

file level, 92
list
 adding users and groups, 93
 removing users and groups, 93
 modifying, 93
 resetting, 94
physical disk resources, 175, 192
planning
 cluster, 175
 network, 176
 protocol, 177
 storage, 176
prerequisites, 13
print server role, removing, 121
print server, configuring, 119
print services, 119
 for UNIX, 123
printer spooler, creating in a cluster, 203
printer, adding, 121
private network adapter, configuring, 181
protocols
 NFS properties settings, 138
 non cluster aware, 190
 parameter settings, 107
 planning, 177
 planning for compatibility, 100
 supported, 19, 106
public network adapter, configuring
 configuring
 public network adapter, 182

Q

quorum disk
 defined, 173
 recommendations, 182

R

rack stability, warning, 14
rapid startup wizard
 defined, 18
redundancy, 18
related documentation, 13
remote access
 iLO port, 166
 Insight Manager, 169
 methods listed, 165
 Remote Desktop, 165
 Telnet Server, 166
 WebUI, 165
Remote Desktop
 defined, 26
 described, 165
 exiting, 27, 154
 improper closure, 27
 opening, 26
 using, 154
remote office deployment, 19
resources, cluster, 172

restarting the server, 24

S

SAN Connection Guide, 178

scheduled shutdown, 24

security

- auditing, 95

- file level permissions, 92

- ownership of files, 97

Server for NFS

- components, 125

- described, 125

services for AppleTalk, installing, 105

services for Macintosh, installing, 105

Services for NFS

- commands, 154

- described, 125

- event logging, 128

setup

- completing, 29

- e-mail alerts, 28

- Ethernet NIC teams, 30, 205

shadow copies

- accessing, 65

- backups, 77

- cache file, 67

- client access, 72

- creating, 68

- defragmentation, 64

- deleting schedule, 69

- described, 61

- disabling, 71

- enabling, 68

- file or folder recovery, 75

- in a cluster, 202

- managing, 65

- mounted drives, 64

- on NFS shares, 74

- on SMB shares, 73

- planning, 61

- properties, viewing, 70

- scheduling, 69

- storage server desktop, 72

- uses, 61

- viewing list, 69

shared disk requirements, 180

shared disks, configuring, 183

shares

- administrative, 100

- creating new, 90, 101

- creating new NCP, 162, 164

- deleting, 102

- managing, 98

- managing for a volume or folder, 91

- modifying NFS properties, 135

- modifying properties, 102

- NCP, 161

- NFS tests, 153

- NFS, creating, 133

NFS, deleting, 135

path, 91

setting up AppleTalk, 106

standard, 100

UNIX, 103

web (HTTP), 104

Windows tab, 102

shutting down the server, 24

simple mapping, 148

simple mappings, 145

smart switch, 210

software

- installing Authentication, 130

software features, 17

software updates, 186

squashed mappings, 145

squashing, 126

Storage Manager, uninstalling, 179

storage reports, 118

storage server

- defined, 17

- desktop, 21

- restarting, 24

- shutting down, 24

- utilities, 17

storage, adding to a cluster, 191

subfolder, navigating to, 88

switch-assisted load balancing, 211

symbols in text, 13

system date, changing, 24

system storage

- managing, 29

system time, changing, 24

T

TCP/IP, configuring on NIC team, 212

technical support, HP, 14

Telnet Server

- enabling, 166

- sessions information, 166

text symbols, 13

time, system, changing, 24

transmit load balancing, 211

U

UNIX

- converting ACL, 146

- group ID, 126

- permissions, 132

- print services, 123

- sharing, 103

- user ID, 126

user access, authenticating, 126

user account, setting up, 182

user credentials, 126

user interfaces, 20

user permissions for NFS, 126

users

adding to permission list, 93

local

 adding, 81

 deleting, 82

 modifying properties, 82

names, managing, 79

NetWare

 adding, 160

 enabling, 160

V

verifying

 connectivity, 182

 disk access, 183

 domain membership, 182

 name resolution, 182

virtual server, 175

virtual server, defined, 172

Volume Shadow Copy Service, 61

volumes

creating new share, 90

creating Novell, 157

managing shares for, 91

navigating to, 87

NCP, 161

W

warning

 rack stability, 14

Web Jetadmin, 123

web sharing, 104

web sites

 HP storage, 15

WebUI

 accessing, 20

 defined, 18

 launching, 165

Windows

 sharing, 102

workgroup environment, 20